Jamming Bandits

SaiDhiraj Amuru[†], Cem Tekin[‡], Mihaela van der Schaar[‡], R. Michael Buehrer[†]

[†]Bradley Department of Electrical and Computer Engineering, Virginia Tech

[‡] Department of Electrical Engineering, UCLA

Email: {adhiraj, rbuehrer}@vt.edu, cmtkn@ucla.edu, mihaela@ee.ucla.edu

arXiv:1411.3652v1 [cs.IT] 13 Nov 2014

Abstract

Can an intelligent jammer learn and adapt to unknown environments in an electronic warfare-type scenario? In this paper, we answer this question in the positive, by developing a cognitive jammer that adaptively and optimally disrupts the communication between a victim transmitter-receiver pair. We formalize the problem using a novel multi-armed bandit framework where the jammer can choose various physical layer parameters such as the signaling scheme, power level and the on-off/pulsing duration in an attempt to obtain power efficient jamming strategies. We first present novel online learning algorithms to maximize the jamming efficacy against static transmitter-receiver pairs and prove that our learning algorithm converges to the optimal (in terms of the error rate inflicted at the victim and the energy used) jamming strategy. Even more importantly, we prove that the rate of convergence to the optimal jamming strategy is sub-linear, i.e. the learning is fast in comparison to existing reinforcement learning algorithms, which is particularly important in dynamically changing wireless environments. Also, we characterize the performance of the proposed bandit-based learning algorithm against multiple static and adaptive transmitter-receiver pairs.

I. Introduction

The inherent openness of the wireless medium makes it susceptible to adversarial attacks. The vulnerabilities of a wireless system can be largely classified based on the capability of an adversary- a) an eavesdropping attack in which the eavesdropper (passive adversary) can listen to the wireless channel and try to infer information (which if leaked may severely compromise data integrity) [1], [2], b) a jamming attack, in which the jammer (active adversary) can transmit energy or information in order to disrupt reliable data transmission or reception [4]-[6] and c) a hybrid attack in which the adversary can either passively eavesdrop or actively jam any ongoing transmission [7], [8]. In this paper, we study the ability of an agent to learn efficient jamming attacks against static and adaptive victim transmitter-receiver pairs.

Jamming has traditionally been studied by using either optimization or game-theoretic or information theoretic principles, see [9]-[15] and references therein. The major disadvantage of these studies is that they assume the jammer has a lot of *a priori* information about the strategies used by the (victim) transmitter-receiver pairs, channel gains, etc., which may not be available in practical scenarios. For

instance, in our prior work [11], we showed that it is not always optimal (in terms of the error rate) to match the jammer's signal to the victim's signaling scheme and that the optimal jamming signal follows a pulsed-jamming strategy. However, these optimal jamming strategies were obtained by assuming that the jammer has *a priori* knowledge regarding the transmission strategy of the victim transmitter-receiver pair. In contrast to prior work (both ours and others), in this paper we develop online learning algorithms that learn the optimal jamming strategy by repeatedly interacting with the victim transmitter-receiver pair. Essentially, the jammer must learn to act in an unknown environment in order to maximize its total reward (e.g., jamming success rate).

Numerous approaches have been proposed to learn how to act in unknown communication environments. A canonical example is reinforcement learning (RL) [16]-[25], in which a radio (agent) learns and adapts its transmission strategy using the transmission success feedback of the transmission actions it has used in the past. Specifically, it learns the optimal strategy by repeatedly interacting with the environment (for example, the wireless channel). During these interactions, the agent receives feedback indicating whether the actions performed were good or bad. The performance of the action taken is measured as a reward or cost, whose meaning and value depends on the specific application under consideration. For instance, the reward can be throughput, the negative of the energy cost, or a function of both these variables. In [18]-[20], Q-Learning based algorithms were proposed to address jamming and anti-jamming strategies against adaptive opponents in multi-channel scenarios. It is well-known that such learning algorithms can guarantee optimality only asymptotically, for example as the number of packet transmissions goes to infinity. However, strategies with only asymptotic guarantees cannot be relied upon in mission-critical applications, where failure to achieve the required performance level will have severe consequences. For example, in jamming applications, the jammer needs to learn and adapt its strategy against its opponent in a timely manner. Hence, the rate of learning matters.

As discussed above, none of the previous works considered the learning performance of physical layer jamming strategies in electronic warfare environments where the jammer has limited to no knowledge about the victim transmitter-receiver pair. Further, the existing learning algorithms [18]-[28] cannot be applied to the problem under consideration because a) none of the existing learning algorithms consider learning over a mixed (mixture of discrete/finite and continuous/infinite actions) action space and b) they do not give any performance guarantees for the jammer's actions. To fill this gap, in this paper, we present novel multi-armed bandit (MAB) algorithms to enable the jammer to learn the optimal physical layer jamming strategies, that were obtained in [11], when the jammer has limited knowledge about the victim. While MAB algorithms have been used in the context of wireless communications to address the selection

of a wireless channel in either cognitive radio networks [21]-[23] or in the presence of an adversary [24], or antenna selection in MIMO systems [25], these works only consider learning over a finite action set. In contrast, the proposed algorithms in this paper enable the jammer to learn the optimal jamming strategies against both static and adaptive victim transmitter-receiver pairs by simultaneously choosing actions from both finite and infinite arm sets (i.e., they can either come from a continuous or a discrete space), that are defined based on the physical layer parameters of the jamming signal. In addition, our algorithms also provide time-dependent (not asymptotic) performance bounds on the jamming performance against static and adaptive victim transmitter-receiver pairs. We note that the algorithms proposed in this paper are novel even within the large area of multi-armed bandits. The major differences between our work and the prior work on multi-armed bandit problems (general works that are not related to jamming) are summarized in Table I.

TABLE I. COMPARISON BETWEEN RELATED BANDIT WORKS

	Finite armed bandits [26]	Continuum armed bandits [28]	Adversarial bandits [27]	Our work
Regret bounds (function of time)	Logarithmic	Sublinear	Sublinear	Sublinear
Action rewards	i.i.d.	i.i.d.	adversarial (worst-case)	i.i.d.
Action set	finite	continuous	finite	mixed

We measure the jamming performance of a learning algorithm using the notion of *regret*, which is defined as the difference between the cumulative reward of the optimal (for example, a strategy that minimizes the throughput of the victim while using minimum energy) jamming strategy when there is complete knowledge about the victim transmitter-receiver pair, and the cumulative reward achieved by the proposed learning algorithm. Any algorithm with regret that scales sub-linearly in time, will converge to the optimal strategy in terms of the average reward. These regret bounds can also provide a rate on how fast the jammer converges to the optimal strategy without having any *a priori* knowledge about the victim's strategy and the wireless channel. As will be discussed in detail later, the feedback considered in this work is also minimal in comparison to the earlier jamming literature [9]-[15].

The rest of the paper is organized as follows. We introduce the system model in Section II. The jamming performance against static and adaptive transmitter-receiver pairs is considered in Sections III and IV respectively, where we develop novel learning algorithms for the jammer and present high confidence bounds for its learning performance. Numerical results are presented in Section V where we discuss the learning behavior in both single and multi-user scenarios and finally conclude the paper in Section VI.

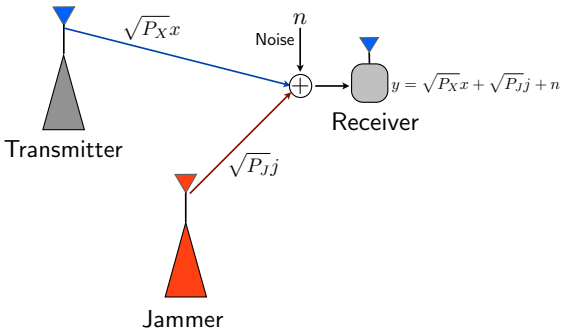

Fig. 1. A wireless environment with victim transmitter-receiver pair and a jammer that intends to disrupt their communication.

II. System Model

We first consider a single jammer and a single victim transmitter-receiver pair in a discrete time setting ($t = 1, 2, \ldots$). We assume that the data conveyed between the transmitter-receiver pair is mapped onto an unknown digital amplitude-phase constellation. The low pass equivalent of this signal is represented as $x(t) = \sum_{m=-\infty}^{\infty} \sqrt{P_x} x_m g(t - mT)$, where P_x is the average received signal power, $g(t)$ is the real valued pulse shape and T is the symbol interval. The random variables x_m denote the modulated symbols assumed to be uniformly distributed among all possible constellation points. Without loss of generality, the average energy of $g(t)$ and modulated symbols $E(|x_m|^2)$ are normalized to unity.[1]

It is assumed that $x(t)$ passes through an AWGN channel (received power is constant over the observation interval) while being attacked by a jamming signal represented as $j(t) = \sum_{m=-\infty}^{\infty} \sqrt{P_J} j_m g(t - mT)$, where P_J is the average jamming signal power as seen at the victim receiver and j_m denote the jamming signals with $E(|j_m|^2) \leq 1$. Assuming a coherent receiver and perfect synchronization, the received signal after matched filtering and sampling at the symbol intervals is given by $y_k = y(t = kT) = \sqrt{P_x} x_k + \sqrt{P_J} j_k + n_k, \ k = 1, 2, ..$ (as shown in Fig. 1), where n_k is the zero-mean additive white Gaussian noise with variance denoted by σ^2. Let $\text{SNR} = \frac{P_x}{\sigma^2}$ and $\text{JNR} = \frac{P_J}{\sigma^2}$. From [11], the optimal jamming signal shares time between two different power levels one of which is 0 and is hence defined by the on-off/pulsing duration ρ. In other words, the jammer sends the jamming signal $j(t)$ at power level JNR/ρ with probability ρ and at power level 0 (i.e., no jamming signal is sent) with probability $1 - \rho$. For more details on the structure of the jamming signals, please see [11]. While the analysis shown in Sections III and IV assumes coherent reception at the victim receiver (i.e., the jamming signal is coherently received along with the transmitter's signal), we consider the effects of a phase offset between these two signals in Section V. The effects of a timing offset between x and j can also be addressed along similar lines, but is skipped in this paper due to a lack of space.

[1] Any signal which follows a wireless standard (such as LTE) would have known parameters such as $g(t)$ and T [29].

III. JAMMING AGAINST A STATIC TRANSMITTER-RECEIVER PAIR

In this section, we consider scenarios where the victim uses a fixed modulation scheme with a fixed SNR. We propose an online learning algorithm for the jammer which learns the optimal power efficient jamming strategy over time, without knowing the victim's transmission strategy.

A. Set of actions for the jammer

At each time t the jammer chooses its signaling scheme, power level and on-off/pulsing duration. A joint selection of these is also referred to as an action. We assume that the set of signaling schemes has N_{mod} elements and the average power level belongs to the set $\mathrm{JNR} \in [\mathrm{JNR}_{\min}, \mathrm{JNR}_{\max}]$.[2] The jamming signal $j(t)$ is defined by the signaling scheme (for example AWGN, BPSK or QPSK) and power level selected at time t. It is shown in [11] that the optimal jamming signal does not have a fixed power level, but instead it should alternate between two different power levels one of which is 0. In other words, the jammer sends the jamming signal j at power level JNR/ρ with probability ρ and at 0 (i.e., no jamming signal is sent) with probability $1-\rho$. Notice that such pulsed-jamming strategies enable the jammer to cause errors with a low average energy but a high instantaneous energy [11]. Therefore, the optimal jamming signal is characterized by the signaling scheme, the average power level and the pulse duration $\rho \in (0, 1]$ which indicates the fraction of time that the jammer is transmitting. The jammer should learn these optimal physical layer parameters by first transmitting the jamming signal and then by observing the reward obtained for its actions.

We formulate this learning problem as a *mixed multi-armed bandit* (mixed-MAB) problem. In contrast to prior work on MAB problems, in a mixed-MAB the action space consists of both finite (signaling set) and continuum (power level, pulse duration) sets of actions. Next, we propose an online learning algorithm called *Jamming Bandits* (JB) where the jammer learns by repeatedly interacting with the transmitter-receiver pair. The jammer receives feedback about its actions by observing the acknowledgment /no acknowledgement (ACK/NACK) packets that are exchanged between the transmitter-receiver pair [31]. The average number of NACKs gives an estimate of the PER which can be used to estimate the SER as as $1-(1-PER)^{1/N_{sym}}$ where N_{sym} is the number of symbols in one packet (other metrics such as throughput or goodput allowed can also be considered [30]). Remember that the SER and PER are

[2]Although we use the variable JNR throughout this paper, it is crucial to notice that the proposed algorithms only need the knowledge of the power with which $j(t)$ is transmitted by the jammer and **do not** need to know the power of the jamming signal as seen at the victim receiver (which depends on the wireless channel whose knowledge is not available to the jammer). There is an unknown but consistent mapping between the jammer's transmit power and JNR. The notation JNR is only used to make the exposition of the Theorems and the algorithms in this paper easier.

functions of the jammer's actions i.e., the signaling scheme, power level and pulse jamming ratio [11] and thereby allow the jammer to learn about its actions.

B. MAB formulation

The actions (also called the *arms*) of the mixed MAB are defined by the triplet [Signaling scheme, JNR, ρ]. The strategy set $\mathcal{S}$, that constitutes JNR and ρ, is a compact subset of $(\mathbb{R}^+)^2$. For each time $t \in \{1, 2, 3, \ldots, n\}$, a cost (or objective) function (feedback metric) $C_t : \{\boldsymbol{\mathcal{J}}, \mathcal{S}\} \rightarrow \mathbb{R}$ is evaluated by the jammer, where $\boldsymbol{\mathcal{J}}$ indicates the set of signaling schemes. Since we are interested in finding power efficient jamming strategies that maximize the error rate at the victim receiver, we define $C_t = \max(SER_t - SER_{target}, 0)/\mathrm{JNR}_t$ or $\max(PER_t - PER_{target}, 0)/\mathrm{JNR}_t$ where JNR_t indicates the average JNR used by the jammer at time t and SER_t, PER_t are the average symbol/packet error rate obtained by using a particular strategy $\{\mathcal{J} \in \boldsymbol{\mathcal{J}}, \mathbf{s} \in \mathcal{S}\}$ at time t and SER_{target}, PER_{target} are the target error rates that should be achieved by the jammer (achieving a target PER is a common constraint in practical wireless systems [29] and this target is defined *a priori*). The dependence of the cost function on the actions taken is unknown to the jammer *a priori* because it is not aware of a) the victim's transmission strategy, b) the power of the signals x and j at the receiver (the probability of error is a function of these parameters as discussed in [11]) and hence needs to be learned over time in order to optimize the jamming strategy. The jammer does this by trying to maximize C_t as it intends to maximize the error rate at the victim receiver using minimum energy.

When the action set is a continuum of arms, most existing MAB works [28] assume that the arms that are close to each other (in terms of the Euclidean distance), yield similar expected costs. Such assumptions on the cost function will at least help in learning strategies that are close to the optimal strategy (in terms of the achievable cost function) if not the optimal strategy [28]. In this paper, for the first time in a wireless communication setting, we prove that this condition indeed holds true i.e., it is not an assumption but rather an intrinsic (proven) feature of our problem and we show how to evaluate the Hölder continuity parameters for these cost functions. Specifically, Theorem 1 shows that this similarity condition indeed holds true when the cost function is SER and extends it to other commonly used cost functions in wireless scenarios. The result in this Theorem is crucial for deriving the regret and high confidence bounds of the proposed learning algorithm.

Formally, the expected or average cost function $\bar{C}(\mathcal{J}, \mathbf{s}) : \{\boldsymbol{\mathcal{J}}, \mathcal{S}\} \rightarrow \mathbb{R}$ is shown to be uniformly locally Hölder continuous with constant $L \in [0, \infty)$, exponent $\alpha \in (0, 1]$ and restriction $\delta > 0$. More specifically, the uniformly locally Hölder continuity condition (described with respect to the continuous

arm parameters) is given by,

$$|\bar{C}(\mathcal{J},\mathbf{s})-\bar{C}(\mathcal{J},\mathbf{s}')|\leq L||\mathbf{s}-\mathbf{s}'||^{\alpha}, \tag{1}$$

for all $\mathbf{s},\mathbf{s}'\in\mathcal{S}$ with $0\leq||\mathbf{s}-\mathbf{s}'||\leq\delta$ [32] ($||\mathbf{s}||$ denotes the Euclidean norm of the continuous 2×1 action vector $\mathbf{s}$). The best strategy $\mathbf{s}^*$ satisfies $\arg\min_{\mathbf{s}\in\mathcal{S}}\bar{C}(\mathcal{J},\mathbf{s})$ for a signaling scheme $\mathcal{J}$. As we will shown next, the algorithms proposed in this paper only require the jammer to know a bound on L and α, since it is not always possible to be aware of the cost function (its dependence on the actions taken) *a priori*.

Theorem 1. *For any set of strategies used by the victim and the jammer, the resultant SER is uniformly locally Hölder continuous.*

Proof: See Appendix A. In an online setting, the Hölder continuity parameters L and α can be estimated if the jammer has knowledge about the victim's transmission strategy, else a bound on L and α works.

We now give an illustrative example for Theorem 1. Consider the scenario where both the jammer and the victim use BPSK modulated signals. The average SER (first we show for the case when $\rho=1$ which will be used to prove the result for $\rho\in(0,1]$) is given by [11]

$$p_e(\mathrm{SNR},\mathrm{JNR})=\frac{1}{4}\left(erfc\left(\frac{\sqrt{\mathrm{SNR}}+\sqrt{\mathrm{JNR}}}{\sqrt{2}}\right)+erfc\left(\frac{\sqrt{\mathrm{SNR}}-\sqrt{\mathrm{JNR}}}{\sqrt{2}}\right)\right), \tag{2}$$

where $erfc$ is the complementary error function. To show the Hölder continuity of the above expression, consider JNR_1 and JNR_2 such that $|\mathrm{JNR}_1-\mathrm{JNR}_2|\leq\delta$, for some $\delta>0$ (i.e., to consider the case of local Hölder continuity). Then by using the Taylor series expansion of the $erfc$ function and ignoring the higher order terms i.e., $erfc(x)\approx1-\frac{2}{\sqrt{\pi}}x+\frac{2}{3\sqrt{\pi}}x^3$, we have

$$p_e(\mathrm{SNR},\mathrm{JNR}_1)-p_e(\mathrm{SNR},\mathrm{JNR}_2)\approx\sqrt{\frac{\mathrm{SNR}}{8\pi}}(\mathrm{JNR}_1-\mathrm{JNR}_2)\leq\sqrt{\frac{\mathrm{SNR}_{\max}}{8\pi}}(\mathrm{JNR}_1-\mathrm{JNR}_2)\,, \tag{3}$$

where $\mathrm{SNR}_{\max}$ relates to the maximum received power level of the victim signal (practical wireless communication devices have limitations on the maximum power levels that can be used). This shows that SER satisfies the Hölder continuity property when $\rho=1$.

For the case of a pulsed jamming signal i.e., $\rho\in(0,1]$, the SER is given by $\rho p_e(\mathrm{SNR},\mathrm{JNR}/\rho)+(1-\rho)p_e(\mathrm{SNR},0)$. The second term is obviously Hölder continuous with respect to the strategy vector $\mathbf{s}=\{\mathrm{JNR},\rho\}$ for $L_1=1,\alpha_1=1$. For the first term, consider the probability of error at the strategies $\mathbf{s}_1=\{\mathrm{JNR}_1,\rho_1\}$ and $\mathbf{s}_2=\{\mathrm{JNR}_2,\rho_2\}$. To prove the Hölder continuity, we consider the expression $\rho_1p_e(\mathrm{SNR},\mathrm{JNR}_1/\rho_1)-\rho_2p_e(\mathrm{SNR},\mathrm{JNR}_2/\rho_2)=\left\{\rho_1p_e\left(\mathrm{SNR},\frac{\mathrm{JNR}_1}{\rho_1}\right)-\rho_1p_e\left(\mathrm{SNR},\frac{\mathrm{JNR}_2}{\rho_1}\right)\right\}+\left\{\rho_1p_e\left(\mathrm{SNR},\frac{\mathrm{JNR}_2}{\rho_1}\right)-\rho_2p_e\left(\mathrm{SNR},\frac{\mathrm{JNR}_2}{\rho_2}\right)\right\}$. Again, the first term in this expression is Hölder continuous with $L_2=\sqrt{\frac{\mathrm{SNR}_{\max}}{8\pi}},\alpha_2=1$ which follows from (3). Using the Taylor series for $erfc$ and after some

manipulations, the second term in this expression can be written as

$$\rho_1 p_e(\text{SNR}, \frac{\text{JNR}_2}{\rho_1}) - \rho_2 p_e(\text{SNR}, \frac{\text{JNR}_2}{\rho_2}) \leq (\rho_1 - \rho_2)\frac{erfc(\text{SNR})}{2}$$
$$\leq \frac{erfc(\text{SNR})}{2}\sqrt{(\text{JNR}_1 - \text{JNR}_2)^2 + (\rho_1 - \rho_2)^2} \triangleq L_3||\mathbf{s} - \mathbf{s}'||^{\alpha_3}. \quad (4)$$

Overall, with $L = 3\min(L_1, L_2, L_3)$ and $\alpha = 1$, the SER obtained under pulsed jamming is also Hölder continuous. In general, since the jammer does not know the victim signals' parameters, it is not aware of the exact structure of the SER expression and hence it can use the worst case L and α (across all possible scenarios that may occur in a real time scenario) to account for the Hölder continuity of C_t.

Corollary 1. *PER and $\max(PER - PER_{target}, 0)/\text{JNR}$ are Hölder continuous.*

Proof: PER can be expressed in terms of the SER. For example, $PER = 1 - (1 - SER)^{N_{sym}}$ when a packet is said to be in error if at least one symbol in the packet is received incorrectly. Since Theorem 1 shows that SER is Hölder continuous, it follows that PER and as a consequence $\max(PER - PER_{target}, 0)/\text{JNR}$ are also Hölder continuous (remember that $\text{JNR} \in [\text{JNR}_{\min}, \text{JNR}_{\max}]$). It is worth noticing that the Hölder continuity parameters L and α depend on the physical layer signaling parameters such as a) the modulation schemes used by the victim and the jammer and b) SNR of the victim signal.

C. Proposed Algorithm

The proposed Jamming Bandits (JB) algorithm is shown in Algorithm 1. At each time t, JB forms an estimate $\hat{C}_t$ on the cost function $\bar{C}$, which is an average of the costs observed over the first $t-1$ time slots. Since some dimensions of the joint action set are continuous, and have infinitely many elements, it is not possible to learn the cost function for each of these values, because it will require a certain amount of time to explore each action from these infinite sets, which thereby cannot be completed in finite time. To overcome this, JB discretizes them and then approximately learns the cost function among these discretized versions. For example, ρ is discretized as $\{1/M, 2/M, \ldots, 1\}$ and JNR is discretized as $\text{JNR}_{\min} + (\text{JNR}_{\max} - \text{JNR}_{\min}) * \{1/M, 2/M, \ldots, 1\}$, where M is the *discretization* parameter. The performance of JB will depend on M, hence, we will also compute the optimal value of M in the following sections.

JB divides the entire time horizon n into several rounds with different durations. Within every round (the duration T of each round is also adaptive as shown in Alg. 1), JB uses a different discretization parameter M to create the discretized joint action set, and learns the best jamming strategy over this set, as shown in Fig. 2. The discretization M increases with the number of rounds. Its value given in line 2 of Algorithm 1 balances the loss incurred due to exploring actions in the discretized set and the loss

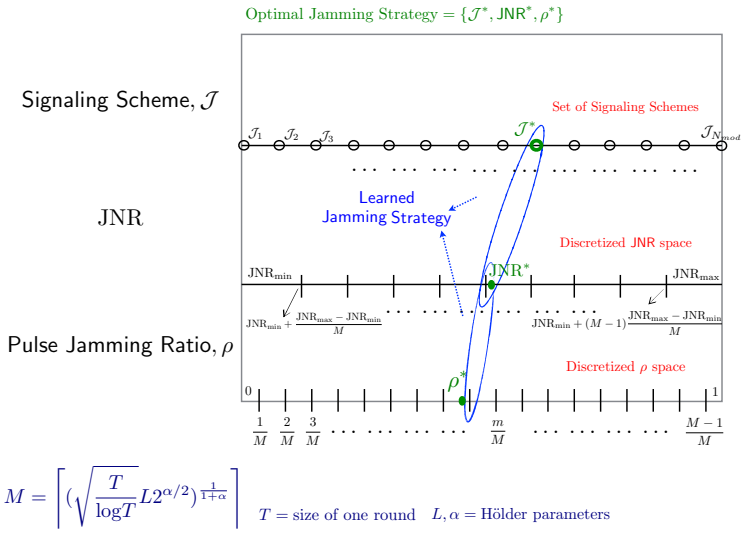
Optimal Jamming Strategy = $\{\mathcal{J}^*, \text{JNR}^*, \rho^*\}$
Signaling Scheme, $\mathcal{J}$
JNR
Pulse Jamming Ratio, ρ
Set of Signaling Schemes
Learned Jamming Strategy
Discretized JNR space
Discretized ρ space
$M = \left\lceil (\sqrt{\frac{T}{\log T}} L 2^{\alpha/2})^{\frac{1}{1+\alpha}} \right\rceil$
T = size of one round
L, α = Hölder parameters

Fig. 2. An illustration of learning in one round of JB. It is possible that the optimal strategy denoted by $\{\mathcal{J}^*, \text{JNR}^*, \rho^*\}$ lies out of the set of discretized strategies. In such a case the jammer learns the best discretized strategy, but based on the value of the discretization parameter M, the loss incurred by using this strategy with respect to the optimal strategy can be bounded using the Hölder continuity condition. The value of the discretization M is shown in the figure and Alg. 1.

incurred due to the sub-optimality resulting from the discretization. The various losses incurred and the derivation of the optimal value for M will be explained in detail in Theorem 2.

Algorithm 1 Jamming Bandits (JB)

T$\leftarrow 1$
while $T \leq n$ **do**
  $M \leftarrow \lceil (\sqrt{\frac{T}{\log T}} L 2^{\alpha/2})^{\frac{1}{1+\alpha}} \rceil$
  Initialize UCB1 algorithm [26] with strategy set $\{\text{AWGN,BPSK,QPSK}\} \times \{1/M, 2/M, \ldots, 1\} \times \text{JNR}_{\min} + (\text{JNR}_{\max} - \text{JNR}_{\min}) * \{1/M, 2/M, \ldots, 1\}$, where $\times$ indicates the Cartesian product.
  for $t = T, T+1, \ldots, \min(2T-1, n)$ **do**
    Choose arm $\{\mathcal{J}_t, \mathbf{s}_t\}$ from UCB1 [26]
    Play $\{\mathcal{J}_t, \mathbf{s}_t\}$ and then estimate $C_t(\mathcal{J}_t, \mathbf{s}_t)$ using the ACK/NACK packets
    For each arm in the strategy set, update its index using $C_t(\mathcal{J}_t, \mathbf{s}_t)$.
  end for
  $T \leftarrow 2T$
end while

Another advantage of JB is that the jammer does not need to know the time horizon n. Time horizon n is only given as an input to JB to indicate the stopping time. All our results in this paper hold true for any time horizon n. This is achieved by increasing the time duration of the inner loop in JB to $2T$ at the end of every round (popularly known as the doubling trick [28]). The inner loop can use any of the standard finite-armed MAB algorithms such as UCB1 [26], which is shown in Algorithm 2 for completeness.

Algorithm 2 Upper confidence bound-based MAB algorithm - UCB1

Initialization: Play each arm once

Loop:

Use signaling scheme $\mathcal{J}$, power JNR, pulse jamming ratio ρ, which maximizes $\hat{C}(\mathcal{J}, \underbrace{\text{JNR}}_{\mathbf{s}}, \rho) + \sqrt{\frac{2\log t}{u_{\mathcal{J},\mathbf{s}}}}$ where t is the time duration since the start of the algorithm, $u_{\mathcal{J},\mathbf{s}}$ is the number of times the arm $\{\mathcal{J}, \mathbf{s}\}$ has been played and $\hat{C}(\mathcal{J}, \underbrace{\text{JNR}}_{\mathbf{s}}, \rho)$ is the estimated average reward obtained from this arm.

D. Upper bound on the regret

For the proposed algorithm, the n-step regret R_n is the expected difference in the total cost between the strategies chosen by the proposed algorithm i.e., $\{\mathcal{J}_1, \mathbf{s}_1\}, \{\mathcal{J}_1, \mathbf{s}_2\}, \ldots, \{\mathcal{J}_n, \mathbf{s}_n\}$ and the best strategy $\{\mathcal{J}^*, \mathbf{s}^*\}$. More specifically, we have $R_n = \mathbf{E}\Big[\sum_{t=1}^{n}\Big(C_t(\mathcal{J}^*, \mathbf{s}^*) - C_t(\mathcal{J}_t, \mathbf{s}_t)\Big)\Big]$, where the expectation is over all the possible strategies that can be chosen by the proposed algorithm. Here we present an upper bound on the cumulative regret that is incurred by the jammer when it uses Algorithm 1 to minimize regret or in other words maximize the cost/objective function.

Theorem 2. *The regret of JB is* $\mathcal{O}(N_{mod} n^{\frac{\alpha+2}{2(\alpha+1)}} (\log n)^{\frac{\alpha}{2(\alpha+1)}})$.

Proof: See Appendix B.

Remark 1. *The upper bound on regret increases as* N_{mod} *increases. This is because the jammer now has to spend more time in identifying the optimal jamming signaling scheme. This does not mean that the jammer is doing worse, since as* N_{mod} *increases, the jamming performance of the* benchmark *against which the regret is calculated also gets better. Hence, the jammer will converge to a better strategy, though it learns more slowly. Further, the regret decreases as* α *increases because higher values of* α *indicate that it is easier to separate strategies that are close (in Euclidean distance) to each other.*

Corollary 2. *The average cumulative regret of JB converges to* 0*. Its convergence rate is given as* $\mathcal{O}(n^{\frac{-\alpha}{2(\alpha+1)}} (\log n)^{\frac{\alpha}{2(\alpha+1)}})$.

The average cumulative regret converges to 0 as n increases. These results establish the learning performance i.e., the rate of learning (how fast the regret converges to 0) of JB and indicate the speed at which the jammer learns the optimal jamming strategy using Algorithm 1. Since the proposed algorithms and hence their regret bounds are dependent only on L and α, which are in turn a function of the various signal parameters such as the modulation schemes used by the victim and the jammer, the wireless channel model i.e., AWGN channel, Rayleigh fading channel etc, the proposed algorithms can be extended to a

wide variety of wireless scenarios by only changing these parameters. The exact values of L and α need not be known in these cases (because the jammer may not have complete knowledge of the wireless channel conditions), the worst case L and α (as shown in the BPSK example below Theorem 1) can be used in the proposed JB algorithm.

E. High Confidence Bounds

The confidence bounds provide an *a priori* probabilistic guarantee on the desired level of jamming performance (e.g., SER or PER) that can be achieved at a given time. We first present the one-step confidence bounds i.e., the instantaneous regret and later show the confidence level obtained on the cumulative regret over n time steps.

The sub-optimality gap Δ_i of the ith arm $\{\mathcal{J}^i, \mathbf{s}^i\}$ (recall that $N_{mod}M^2$ arms can be chosen in one round of JB), is defined as $\bar{C}(\mathcal{J}^*, \mathbf{s}^*) - \bar{C}(\mathcal{J}^i, \mathbf{s}^i)$. We say that an arm is sub-optimal if it belongs to the set $U_>$ (set of arms whose sub-optimality gap exceeds a threshold based on the required jamming confidence level) which is defined in detail in Appendix C. Let $u_i(t)$ denote the total number of times the ith arm has been chosen until time t and $U(T)$ indicate the set of time instants $t \in [1, T]$ for which $u_i(t) \leq \frac{8\log(T)}{\Delta_i^2}$ for some sub-optimal arm $i \in \mathcal{U}_>$.

Theorem 3. *(i) Let* $\delta = 2 \times 2^{\frac{3\alpha+2}{2(1+\alpha)}} L^{\frac{1}{1+\alpha}} \left(\frac{\log T}{T}\right)^{\frac{\alpha}{2(1+\alpha)}}$ *and* M *be defined as in Algorithm 1. Then for any* $t \in [1, T]\backslash U(T)$*, with probability at least* $1 - 2(N_{mod} + M^2)t^{-4}$*, the expected cost of the chosen jamming strategy* $(\mathcal{J}_t, \mathbf{s}_t)$ *is at most* $\bar{C}(\mathcal{J}^*, \mathbf{s}^*) + \delta$*. In other words,*
$P\left(\bar{C}(\mathcal{J}^*, \mathbf{s}^*) - \bar{C}(\mathcal{J}_t, \mathbf{s}_t) > \delta\right) \leq 2(N_{mod} + M^2)t^{-4}$.
(ii) We also have

$$\begin{aligned} E[|U(T)|] &\leq \sum_{t=1}^{T} P(\textit{a sub-optimal arm } i \in \mathcal{U}_> \textit{ is chosen at } t) \\ &\leq 8 \sum_{i \in \mathcal{U}_>} \left(\frac{\log T}{\Delta_i^2}\right) + \left(1 + \frac{\pi^2}{3}\right) |U_>|, \end{aligned}$$

which means that our confidence bounds hold in all except logarithmically many time slots in expectation.

Proof: See Appendix C.

Remark 2. *A lower bound on the sub-optimality gap i.e.,* $\Delta_{\min} = \min_{i \in \mathcal{U}_>} \Delta_i$*, can be used to approximately estimate* $U(T)$*. For instance, in a wireless setting when* SER *is used as the cost function, if the jammer is aware of the smallest tolerable error in* SER *that is allowed, then it can approximately evaluate* $U(T)$*. A detailed discussion on how the jammer can estimate* $U(T)$ *is given in Appendix C.*

Corollary 3. *The one-step regret converges to zero in probability i.e.,*

$$\lim_{T\to\infty}\left(\lim_{t\to T} P\left(\bar{C}(\mathcal{J}^*,\mathbf{s}^*)-\bar{C}(\mathcal{J}_t,\mathbf{s}_t)>\delta\right)\right)=0.$$

Theorem 3 can be used to achieve desired confidence levels about the jamming performance, which is particularly important in military settings. In order to achieve a desired confidence level (e.g., about the SER inflicted at the victim receiver) δ at each time step, the probability of choosing a jamming action that incurs regret more than δ must be very small. In order to achieve this objective, the jammer can set M as $\max\{(\frac{2^{\frac{\alpha+4}{2}}L}{\delta})^{1/\alpha}, \lceil(\sqrt{\frac{T}{\log T}}L2^{\alpha/2})^{\frac{1}{1+\alpha}}\rceil\}$. By doing this, the jammer will not only guarantee a small regret at every time step, but also chooses an arm that is within δ of the optimal arm at every time step with high probability. Hence, the one time step confidence about the jamming performance can be translated into overall jamming confidence. It was, however, observed that the proposed algorithm performs significantly better than predicted by this bound (Section V).

Theorem 4. *For any signaling scheme $\mathcal{J}$ chosen by the jammer, $P\Big(\sum_{t=1}^{T}(\bar{C}(\mathcal{J},\mathbf{s}^*)-\bar{C}(\mathcal{J},\mathbf{s}_t))>\left(\frac{8}{3\epsilon}\left(\frac{T}{\log T}\right)^{\frac{4}{1+\alpha}}\right)^{1/3}\Big)<\epsilon,\ \forall\ \epsilon>0$.*

Proof: See Appendix C. Using Theorem 4, a confidence bound on the overall cumulative regret defined as $\sum_{t=1}^{T}[\bar{C}(\mathcal{J}^*,\mathbf{s}^*)-\bar{C}(\mathcal{J}_t,\mathbf{s}_t)]$ can be directly obtained as shown in Appendix C. This bound indicates the overall confidence acquired by the jammer. The regret performance of JB will be discussed in more detail via numerical results in Section V.

Theorem 5. *Let $\delta=2\times 2^{\frac{5\alpha+4}{2(1+\alpha)}}L^{\frac{1}{1+\alpha}}\left(\frac{\log T}{T}\right)^{\frac{\alpha}{2(1+\alpha)}}$ and M be defined as in JB. Then, for any $t\in[1,T]\backslash U(T)$, the jammer knows that with probability at least $1-2(N_{mod}+M^2)t^{-4}-t^{-16}$, the true expected cost of the optimal strategy is at most $\hat{C}(\mathcal{J}_t,\mathbf{s}_t)+\delta$, where $\hat{C}(\mathcal{J}_t,\mathbf{s}_t)$ is the sample mean estimate of $\bar{C}(\mathcal{J}_t,\mathbf{s}_t)$, the expected reward of strategy $(\mathcal{J}_t,\mathbf{s}_t)$ selected by the jammer at time t.*

Proof: See Appendix D. Theorem 5 presents a high confidence bound on the estimated cost function of any strategy used by the jammer. Such high confidence bounds (Theorems 3-5) will enable the jammer to make decisions on the jamming duration and jamming budget, which is explained below with an example. Again, this is a worst case bound and the proposed algorithm performs much better than predicted by the bound as will be discussed in detail in Section V.

Remark 3. *Fig. 3 summarizes the importance and usability of Theorems 3 and 5 in realtime wireless communication environments. The high confidence bounds for the regret help the jammer decide the*

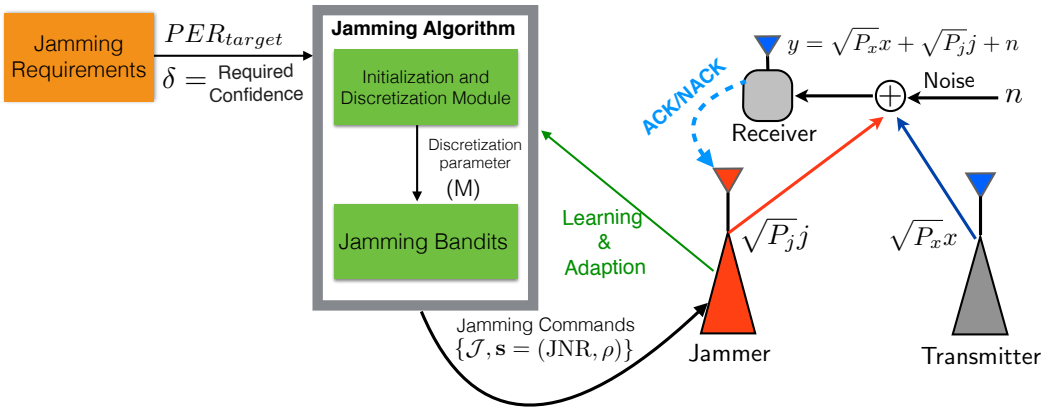

Fig. 3. Using Theorems 3 and 5 in a real time jamming environment.

number of symbols (or packets) to be jammed to disrupt the communication between the victim transmitter-receiver pair. For example, such confidence is necessary in scenarios where the victim uses erasure or rateless codes and/or HARQ-based transmission schemes. In the case of rateless codes, a message of length N is encoded into an infinitely long new message sequence of length $\hat{N} >> N$ (for example, by using random linear combinations) out of which any N are linearly independent. Upon successfully receiving N such messages, the entire message can be recovered. Under such scenarios, the high confidence bounds help the jammer to decide the number of packets/ time instants to jam successfully in order to disrupt the wireless link between the transmitter-receiver pair.

For instance, when $M = 15$, we have at large time t, $\delta > 0.01$, i.e., $P(SER^* - \hat{SER}_t > 0.01) = 0$, where SER^* is the optimal average SER achievable and $\hat{SER}_t$ is the estimated SER achieved by the strategy used at time t. If the jammer estimates SER as 0.065 then the best estimate of the SER^* indicates that it is less than or equal to 0.075. Using such knowledge, the jammer can identify the minimum number of packets it has to jam so as to disrupt the communication and prevent the exchange of a certain number of packets (which in applications such as video transmission can completely break down the system). As an example, consider the case when packets of length 100 symbols are exchanged and that a packet is said to be in error only when there are more than 10 errors in the packet. Thus, in order to jam 100 packets successfully the jammer needs to affect at least 463 packets on an average if SER^* (which corresponds to $PER = 0.2167$) was achievable. However, since it can only achieve $\hat{SER} = 0.065$ i.e., $\hat{PER} = 0.1153$, it has to jam at least 865 packets on an average to have sufficient confidence regarding its jamming performance. The jammer can accordingly plan its energy budget/ jamming duration etc. by using such knowledge.

F. Improving convergence via arm elimination

When the number of signaling schemes that the jammer can choose from is large or when α is small (i.e., it is difficult to separate the arms that are close to each other), then the learning speed using JB can be relatively slow. We now present an algorithm to improve the learning rate and convergence speed of JB under such scenarios. In order to achieve this, Algorithm 1 is modified to use the UCB-Improved algorithm [34] inside the inner loop of JB instead of UCB1. The UCB-Improved algorithm eliminates sub-optimal arms (that are evaluated in terms of the mean rewards and the confidence intervals), in order to avoid exploring the sub-optimal arms (which is important in electronic warfare scenarios). The modified algorithm and the associated UCB-Improved algorithm are shown in Algorithms 3 and 4 respectively.

Algorithm 3 Jamming Bandits with Arm Elimination

T$\leftarrow 1$
while $T \leq n$ **do**
  Initialize UCB-Improved [34] algorithm with the strategy set $\{$AWGN,BPSK,QPSK$\}\times \{1/M, 2/M, \ldots, 1\}\times$JNR$_{\min}$+(JNR$_{\max}$-JNR$_{\min}$)*$\{1/M, 2/M, \ldots, 1\}$, where $\times$ indicates the Cartesian product.
  for $t = T, T+1, \ldots, \min(2T-1, n)$ **do**
    Use the UCB-Improved [34] MAB Algorithm to eliminate sub-optimal arms
  end for
  $T \leftarrow 2T$
end while

Algorithm 4 UCB-Improved

Input the set of arms A and time horizon T
$\tilde{\Delta}_0 = 0, B_0 = A$
for rounds $m = 0, 1, 2, \ldots, \frac{1}{2}\log_2 \frac{T}{e}$ **do**
  Arm Selection
  If $|B_m| > 1$, choose each arm in B_m for $n_m = \lceil \frac{2\log(T\tilde{\Delta}_m^2)}{\tilde{\Delta}_m^2} \rceil$
  Else choose the remaining arm until time T
  Arm Elimination
  Delete arm i in the set B_m for which $\left(\bar{C}_i + \sqrt{\frac{\log(T\tilde{\Delta}_m^2))}{2n_m}}\right) < \max_{j\in B_m}\left(\bar{C}_j - \sqrt{\frac{\log(T\tilde{\Delta}_m^2))}{2n_m}}\right)$ to obtain the set of new arms B_{m+1}; $\bar{C}_i$ is the average cost incurred by playing arm i.
  Reset $\tilde{\Delta}_m : \tilde{\Delta}_{m+1} = \tilde{\Delta}_m/2$.
end for

To obtain the value of M i.e., the discretization for JNR and ρ, we used numerical optimization tools to solve $TL\left(\frac{2}{M^2}\right)^{\frac{\alpha}{2}} - \left(\sqrt{M^2T}\frac{\log(M^2\log(M^2))}{\sqrt{\log(M^2)}}\right) = 0$. See Appendix E for more details. Later in Section V, we show the benefits of using this algorithm via numerical simulations. The regret bounds can be derived along similar lines to Theorems 1-5 by using the properties of the UCB-Improved algorithm [34].

IV. Learning jamming strategies against a time-varying user

In this section, we consider scenarios where the victim transmitter-receiver pair can choose their strategies in a time-varying manner.[3] We specifically consider two scenarios a) when the victim changes its strategies in an i.i.d. fashion and b) when the victim is adapting its transmission strategies to overcome the interference seen in the wireless channel.[4] The worst case jammer's performance can be understood by considering a victim that changes its strategies in an i.i.d. fashion. For example, such i.i.d. strategies are commonly employed in a multichannel wireless system where the victim can randomly hop onto different channels (either in a pre-defined or an un-coordinated fashion [35]) to probabilistically avoid jamming/ interference. The randomized strategies chosen by the victim can confuse the jammer regarding its performance. For instance, if the jammer continues using the same strategy irrespective of the victim's strategy, then the jammers' performance will be easily degraded. However, if the jammer is capable of anticipating such random changes by the victim and learns the jamming strategies, then it can disrupt the communication irrespective of the victims' strategies.

We assume that the victim can modify its power levels and the modulation scheme to adapt to the wireless environment (the most widely used adaption strategy [33]). Again we allow the jammer to learn the optimal jamming strategy by optimizing the 3 actions, namely signaling scheme, JNR and ρ as before. The jammer has to learn its actions without any knowledge regarding the victim's strategy set and any possible distribution that the victim may employ to choose from this strategy set. We use Algorithm 1 and not Algorithm 3 to address such dynamic scenarios because eliminating arms in such a time-varying environment may not always be beneficial. For example, a certain arm might not be good against one strategy used by the victim but might be the optimal strategy when the victim changes its strategy.

While the regret bounds presented below assume that the victim employs a random unknown distribution over its strategy set and chooses its actions in an i.i.d. manner (also referred to as stochastic strategies) i.e., scenario (a) mentioned earlier, we discuss the jammer's performance against any strategy (i.e., without any predefined distribution over the strategies, for example, increase the power levels when the PER increases) employed by the victim (which includes scenario (b)) in Section V.

---

[3]The model considered in this formulation is different from the *adversarial scenarios* studied in the context of MAB algorithms [27]. In the adversarial bandit cases, the adversary (or the victim in this current context) observes the action of the jammer and then assigns a reward function either based on the jammers' current action or on the entire history of jammers' actions. However, in the current scenario we assume that the user picks a strategy in an i.i.d manner independent of the jammer. Considering learning algorithms in adversarial scenarios is reserved for future work.

[4]While the victim is not entirely adaptive against the jammers' strategies, it is adaptive in the sense that it can choose from a set of strategies to overcome the jamming/ interference effects. For example, it can be adaptive based on the PER seen at the victim receiver. This scenario is discussed in detail in Section V.

1) Upper bound on the regret: Let $\{p_i\}_{i=1}^{|\mathcal{P}|}$ denote the probability distribution with which the victim selects its strategies in an i.i.d manner, from a set consisting of $|\mathcal{P}|$ number of possible strategies. The jammer is not aware of this distribution chosen by the victim and needs to learn the optimal strategy by repeatedly interacting with the victim. The regret under such scenarios is defined as $R_n = \mathbf{E}\Big[\sum_{t=1}^{n}\Big(C_t(\mathcal{J}^*,\mathbf{s}^*) - C_t(\mathcal{J}_t,\mathbf{s}_t)\Big)\Big]$, where the expectation is over the random strategies chosen by the jammer as well as the victim (which is different from the formulation in Section III). Thus, the above expression can be re-written as $R_n = \mathbf{E}\Big[\sum_{t=1}^{n}\sum_{i=1}^{|\mathcal{P}|} p_i\Big(C_t^i(\mathcal{J}^*,\mathbf{s}^*) - C_t^i(\mathcal{J}_t,\mathbf{s}_t)\Big)\Big]$, with C_t^i indicating the cost function when the victim uses strategy i with probability p_i and the expectation is now taken only over the strategies chosen by the jammer.

Theorem 6. *The regret of JB when the victim employs stochastic strategies is* $\mathcal{O}(N_{mod} n^{\frac{\alpha+2}{2(\alpha+1)}} (\log n)^{\frac{\alpha}{2(\alpha+1)}})$.

Proof: See Appendix F. This is an upper bound on the cumulative regret incurred by JB under such stochastic scenarios. Similar to the regret incurred by JB in Theorem 1, the regret under stochastic cases also converges to 0 as $\mathcal{O}(n^{\frac{-\alpha}{2(\alpha+1)}} (\log n)^{\frac{\alpha}{2(\alpha+1)}})$. The one step confidence bounds similar to Theorems 3-5 can be derived even in this case but are skipped due to lack of space.

Remark 4. *When the victim is adapting its strategies based on the error rates observed over a given time duration (as is typically done in practical wireless communication systems), we show that by employing sliding-window based algorithms, the jammer can effectively track the changes in the victim and jam it in a power efficient manner. This is discussed more in detail in the next section.*

V. Numerical Results

We first discuss the learning behavior of the jammer against a transmitter-receiver pair that employs a static strategy and later consider the performance against adaptive strategies. To validate the learning performance, we compare the results against the optimal jamming signals that are obtained when the jammer has complete knowledge about the victim [11]. It is assumed that the victim and the jammer send 1 packet with 10000 symbols at any time t. A packet is said to be in error if at least 10% of the symbols are received in error at the victim receiver so as to capture the effect of error correction coding schemes. The minimum and the maximum $\mathrm{SNR}, \mathrm{JNR}$ levels are taken to be $0\ dB$ and $20\ dB$ respectively. The set of signaling schemes for the transmitter-receiver pair is $\{BPSK, QPSK\}$ and for the jammer is $\{AWGN, BPSK, QPSK\}$[5] [11] i.e., $N_{mod} = 3$.

[5]It is very easy to extend the results in this paper and [11] to PAM and QAM signals of any constellation size.

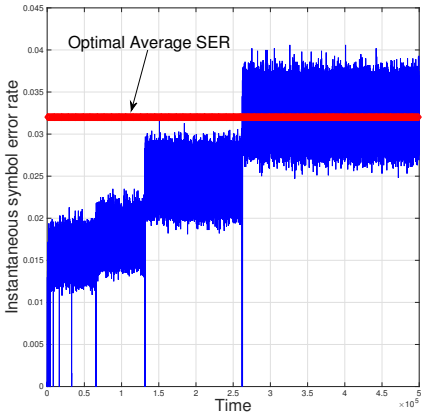

Fig. 4. Instantaneous SER achieved by the JB algorithm when JNR $= 10dB$, SNR $= 20dB$ and the victim uses BPSK.

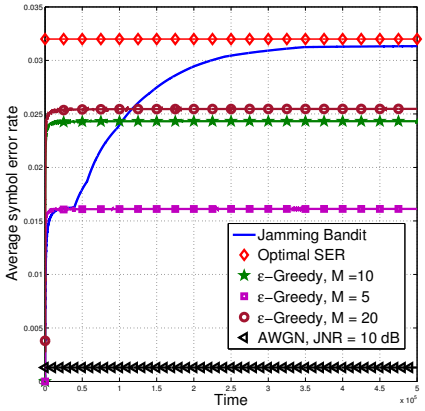

Fig. 5. Average SER achieved by the jammer when JNR $= 10dB$, SNR $= 20dB$ and the victim uses BPSK. The jammer learns to use BPSK with $\rho = 0.078$ using JB. The learning performance of the ϵ-greedy learning algorithm with various discretization factors M is also shown.

A. Fixed user strategy

The jammer uses SER or PER inflicted at the victim receiver (estimated using the ACK and NACK packets) as feedback to learn the optimal jamming strategy. We first consider a scenario where the JNR is fixed and the jammer can optimize its jamming strategy by choosing the optimal signaling scheme $\mathcal{J}^*$ and the associated pulse jamming ratio ρ^*. These results enable comparison with previously known results obtained via an optimization framework with full knowledge about the victim as discussed in [11]. Note that unlike [11], the jammer here does not know the signaling parameters of the victim signal, and hence it cannot solve an optimization problem to find the optimal jamming strategy. In contrast, it learns over time the optimal strategy by simply learning the expected reward of each strategy it tries.

Figs. 4-7 show the results obtained in this setting (fixed SNR, modulation scheme for the victim and fixed JNR). For a fair comparison with [11], we initially assume that the jammer can directly estimate the SER inflicted at the victim receiver. We will shortly discuss the more practical setting in which the jammer can only estimate PER. In all these figures, it is seen that the jammers' performance converges to that of the optimal jamming strategies [11]. For example, in Figs. 4 and 5, when the victim transmitter-receiver pair exchange a BPSK modulated signal at SNR $= 20$ dB, the jammer learns to use BPSK signaling at JNR $= 10$ dB and $\rho = 0.078$ which is in agreement with the results presented in [11].

Fig. 4 shows the instantaneous learning performance of the jammer in terms of the SER achieved by using the JB algorithm. The variation in the achieved SER after convergence is only due to the wireless channel. The time instants at which the SER varies a lot, i.e., the dips in SER seen in these results

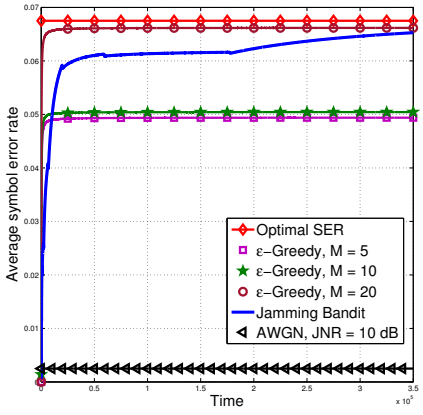

Fig. 6. Learning the optimal jamming strategy when JNR $= 10dB$, SNR $= 20dB$ and the victim uses QPSK modulation scheme. The jammer learns to use QPSK signaling scheme with $\rho = 0.087$.

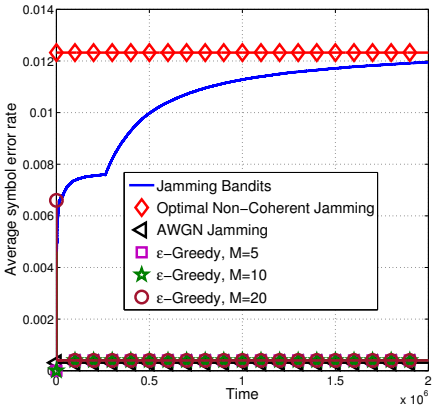

Fig. 7. Average SER achieved by the jammer when JNR $= 10dB$, SNR $= 20dB$ and the victim uses BPSK and there is a phase offset between the two signals. The jammer learns to use BPSK with $\rho = 0.051$ using JB. The learning performance of the ϵ-greedy learning algorithm with various discretization factors M is also shown.

are due to the exploration phases performed when a new value of discretization i.e., M is chosen by the algorithm (recall from Algorithm 1 that for every round the discretization M is re-evaluated). Fig. 5 shows the average SER attained by this learning algorithm. Also shown in Fig. 5 is the performance of the the ϵ-greedy learning algorithm [26] with exponentially decreasing exploration probability $\epsilon(t) = \epsilon^{\frac{t}{10}}$ (initial exploration probability is taken to be 0.9) and various discretization factors M. In the ϵ-greedy learning algorithm, the jammer explores (i.e., it tries new strategies) with probability $\epsilon(t)$ and exploits (i.e., uses the best known strategy that has been tried thus far) with probability $1 - \epsilon(t)$. It is seen that unless the optimal discretization factor M is known (so that the optimal strategy is one among the possible strategies that can be chosen by the ϵ-Greedy algorithm), the ϵ-greedy algorithm performs significantly worse in comparison to the novel bandit-based learning algorithm.

Similar results were observed in the case of QPSK signaling as seen in Fig. 6. Notice that while the ϵ-greedy algorithm with discretization $M = 20$ did not achieve satisfactory results in the BPSK signaling scenario, it achieved close to optimal results in the QPSK signaling scenario as seen in Fig. 6. Thus, the performance of the ϵ-greedy algorithm highly depends on M, and it can be sub-optimal if M is chosen incorrectly. However, in our learning setting it is not possible to know the optimal M *a priori*. Also, the performance of AWGN jamming (which is the most widely used jamming signal [14], [33] when the jammer is not intelligent) is significantly lower than the performance of JB. The algorithms behave along similar lines in a non-coherent scenario where there is a random unknown phase offset between the jamming and the victim signals, as seen in Fig. 7. The jammer learned to use BPSK signaling at

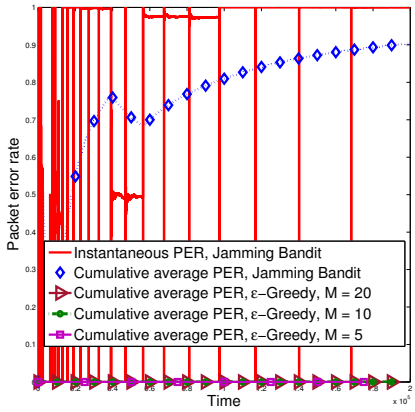
Packet error rate
Time
Instantaneous PER, Jamming Bandit
Cumulative average PER, Jamming Bandit
Cumulative average PER, ε–Greedy, M = 20
Cumulative average PER, ε–Greedy, M = 10
Cumulative average PER, ε–Greedy, M = 5

Fig. 8. Average PER inflicted by the jammer at the victim receiver, SNR $=$ 20 dB, victim uses BPSK and JNR $=$ 10 dB. The jammer learns to use BPSK signaling scheme with $\rho =$ 0.23.

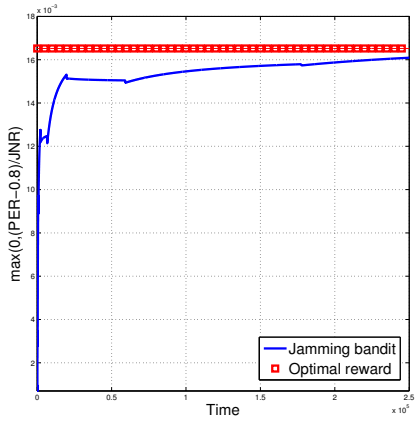

Fig. 9. Average reward obtained by the jammer against a BPSK modulated victim, $SNR = 20$ dB. The optimal reward is obtained via grid search with discretization $M = 100$.

$\rho = 0.051$ while the optimal jamming signal derived in [11] indicates that $\rho^* = 0.06$ when $JNR = 10\ dB$ and $SNR = 20\ dB$.

Now that we have established the performance of the proposed learning algorithm by comparing with previously known results, we now consider the performance of the learning algorithm in terms of the PER which is a more relevant and practical metric to be considered in wireless environments. Further, it is also easy for the jammer to estimate PER by observing the ACKs/NACKs exchanged between victim receiver and transmitter via the feedback channel [31]. Fig. 8 shows the learning performance of various algorithms in terms of the average PER inflicted by the jammer at the victim receiver. While the jammer learns to use BPSK as the optimal signaling scheme, the optimal ρ value learned in this case is 0.23 which is different from the value of ρ learned in Fig. 5. This is because PER is used as the cost function in learning the jamming strategies. It is clear that both the AWGN jamming and ϵ-greedy learning algorithm (that uses a sub-optimal value of M) achieve a $PER = 0$ based on the SER results in Fig. 5. Even in this case, JB outperforms traditional jamming techniques that use AWGN or the ϵ-greedy learning algorithm.

The cost function is taken as $\max(0, (PER(t) - 0.8)/\mathrm{JNR}(t))$ (the cost function remains to be Hölder continuous and is bounded in $[0, 1]$) to ensure that we choose only those strategies which achieve at least 80% PER (remember, the jammer intends to maximize this cost/objective function). Fig. 9 compares the learning performance of JB with respect to the optimal strategy and Fig. 10 shows the confidence levels as predicted by the one-step regret bound in Theorem 3 and that achieved by JB. The optimal reward is estimated by performing an extensive grid search $(M = 100)$ over the entire strategy set. The steps in

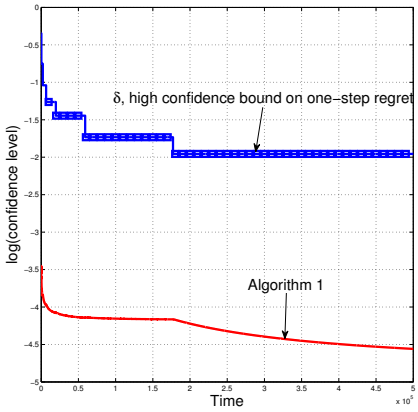

Fig. 10. Confidence level (optimal reward-achieved reward) predicted by Theorem 3 and that achieved by JB.

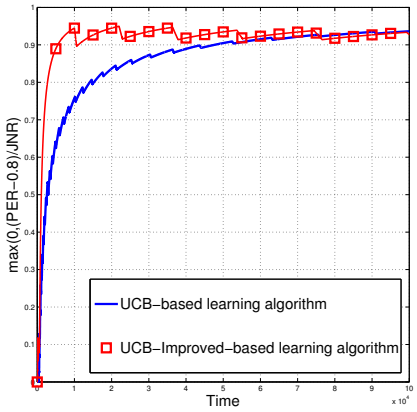

Fig. 11. Learning the jamming strategies by using arm-elimination. The victim uses BPSK with $\mathrm{SNR} = 20$dB. The jammer learned to use BPSK with $\mathrm{JNR} = 15$ dB and $\rho = 0.22$.

$\log\delta$ seen in Fig. 10 are due to change in the discretization M as shown in Algorithm 1. As mentioned before, the algorithm performs much better than predicted by the high confidence bound (evidenced by a lower value of δ).

Fig. 11 shows the learning results obtained by using Algorithm 3 i.e., JB uses the UCB-Improved algorithm in the inner loop instead of the UCB1 algorithm. It shows the learning performance of Algorithms 1 and 3 in one inner loop iteration when $T = 10^5$ (i.e., for one value of discretization M evaluated as shown in Algorithm 1). It is seen that the Algorithm 3 converges faster in comparison to the earlier approach as the algorithm eliminates sub-optimal arms and thereby only exploits the best jamming strategy. Even in this case the jammer learned to use BPSK signaling scheme against a BPSK-modulated victim signal. Further notice that the algorithm converges in about 10000 time steps in this case as opposed to > 50000 time steps using JB. Recall that in the simulations we assume that one packet is sent every time instant and hence in order to obtain reliable estimates of the performance of each jamming strategy, the jammer requires about 10000 time instants.

B. Jamming Performance Against an Adaptive Victim

We first assume that the victim employs a uniform distribution over its strategy set i.e., it chooses uniformly at random (at every time instant) a power level in the range $[\mathrm{SNR}_{\min}, \mathrm{SNR}_{\max}]$ and the modulation scheme from the set $\{$BPSK,QPSK$\}$. The performance of JB when the victim employs such a stochastic strategy is shown in Fig. 12. Again, the superior performance of the bandit-based learning algorithm when compared to the traditionally used AWGN jamming and naive learning algorithms such

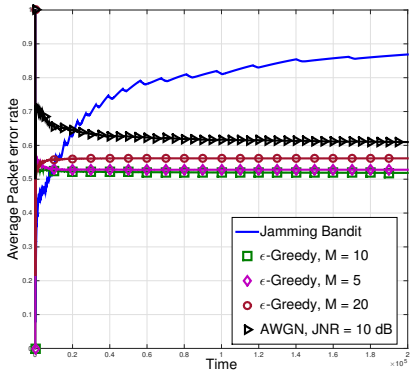

Fig. 12. Learning jammers' strategy against a stochastic user. The victim transmitter-receiver pair use a uniformly random signaling scheme that belongs to the set {BPSK,QPSK} and random power level in the range $[0, 20]$ dB.

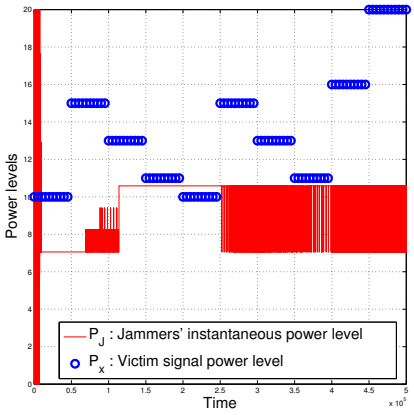

Fig. 13. Learning against a victim with time-varying strategies. The figure shows the power levels adaptation by the jammer and that used by the victim.

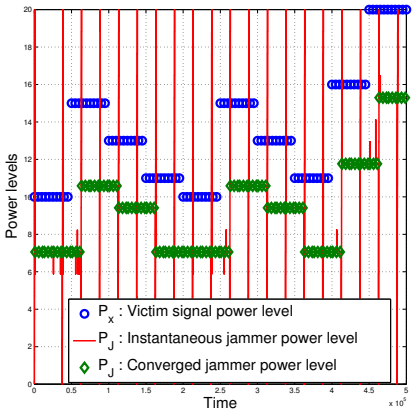

Fig. 14. Learning against a victim with time-varying strategies. The figure shows the power levels adaptation by the jammer using a drifting algorithm and that used by the victim.

as ϵ-Greedy is proved from these results.[6]

When the victim changes its strategy rapidly, JB cannot track the changes perfectly as seen in Fig. 13 because it learns over all past information, and prior information may not convey knowledge about the current strategy used by the victim which can be completely different from the prior strategy. In such cases, it is important to learn only from recent past history, which can be achieved by using JB on a recent window of past history (for instance, a sliding window-based algorithm to track changes in the environment) [36]. Specifically, we use the concept of drifting [36] to adapt to the victim's strategy. In this algorithm, each round i (which is of T time steps, where $T = 2^i$) is divided it into several frames each of W time instants. Within each frame, the first $W/2$ time steps, are termed as the passive slot and the second $W/2$ time instants are termed as the active slot. In the first frame, both the slots will

[6]It is worth noticing that model-free learning algorithms such as Q-Learning and SARSA [20] cannot be employed in the scenarios considered in this paper because it is assumed that the jammer cannot observe any of the environment parameters such as the victim's modulation scheme and power levels. However, it is expected that the performance of the learning algorithms can be improved when such additional information is available, which is typically the case in optimization-based algorithms.

be taken to be active slots. Each passive slot overlaps with the active slot of the previous frame. If time t belongs to active slot of frame w, then actions are taken as per the UCB1 indices evaluated in this particular frame w. However, if it belongs to the passive slot of frame w, which is taken to overlap with the active slot of frame $w-1$, then it takes actions as per the indices of the frame $w-1$, but updates the UCB1 indices so that it can be used in frame w. Specifically, at the start of every frame w, the counters and mean reward estimates are all reset to zero and when actions are taken in the passive slot of frame w, these counters and reward estimated are updated so as to be used in the active slot. Thus when the algorithm enters the active slot of frame w, it already has some observations using which it can exploit without wasting time in the exploration phase. Such splitting of the time horizon will enable the jammer to quickly adapt to the victim's varying strategies. Please see [36] for more details on the drifting algorithm. Specifically, we consider the drifting algorithm with a window length $W = 25000$.

Fig. 14 shows the jammers' power level adaption when the victim is randomly varying its power levels across time and the jammer employs the drifting algorithm in conjunction with JB. The dips seen at regular intervals in Fig. 14 are due to the proposed sliding window-based algorithm where the user resets the algorithm at regular intervals to adapt to the changing wireless environment. The PER achieved by this algorithm is similar to the results shown in Figs. 8, 10 in comparison to other jamming techniques. While Fig. 14 considered the case when the victim changes its power levels randomly, the jammer can also easily track the victim when it employs commonly used adaption strategies such as increasing the power levels when PER increases and vice versa. These results successfully illustrate the adaptive capabilities of the proposed learning algorithms that can overcome the difficulties faced by JB as shown in Fig. 13.

C. Multiple Victims

In this subsection, we consider a case when the jammer uses an omnidirectional antenna and intends to jam two victims in a network. Interesting scenarios arise in this scenario because the jammer has to optimize its jamming strategy based on the PER of both the victims. For example, when both the victims use BPSK, the jammer will learn to use BPSK signaling scheme but the power level at which it should jam depends on the relative power levels of both the victims. Several factors such as path loss, shadowing etc. akin to practical wireless systems can be introduced into this problem, but we are mainly interested in understanding the learning performance of the jammer. Hence we ignore these physical layer parameters and assume that both the victims are affected by the jamming signal with the same JNR. The jammer considers the mean packet error rate seen at both these victims as feedback with target mean $PER = 0.8$, in order to learn the performance of its actions.

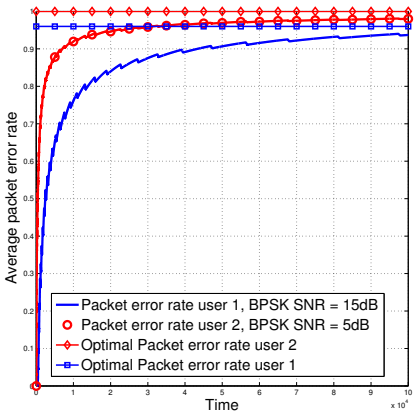
Average packet error rate
Time
Packet error rate user 1, BPSK SNR = 15dB
Packet error rate user 2, BPSK SNR = 5dB
Optimal Packet error rate user 2
Optimal Packet error rate user 1

Fig. 15. PER achieved by the jammer against 2 users, user 1 uses BPSK at 15dB and user 2 sends BPSK at 5dB. The jammer learns to use BPSK signal with power 13dB and $\rho = 0.46$.

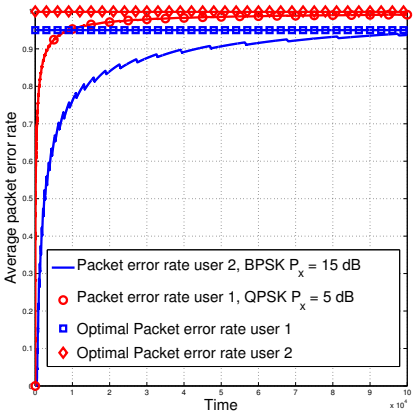
Average packet error rate
Time
Packet error rate user 2, BPSK P_x = 15 dB
Packet error rate user 1, QPSK P_x = 5 dB
Optimal Packet error rate user 1
Optimal Packet error rate user 2

Fig. 16. PER achieved by the jammer against 2 users, user 1 sends QPSK at 5dB and user 2 sends BPSK at 15dB. The jammer learns to use BPSK signal with power 11.25dB and $\rho = 0.25$.

Fig. 15 shows the learning performance of the jammer against 2 users that employ BPSK signaling at different power levels. It is seen that the jammer learns to use BPSK signaling as well (since BPSK is optimal to be used against BPSK signaling as discussed in [11]). Similar learning results were achieved when both the users employ QPSK signaling. Figs. 16 shows the learning performance when one user uses QPSK and and the other user uses BPSK. It was observed that when the victim with BPSK has higher power than QPSK victim, the jammer learns to use the BPSK jamming signal and vice versa. This again agrees with previous results which show that BPSK (QPSK) is better to jam a BPSK (QPSK) signal. Also, the learning algorithm performs comparably well to the optimal strategy obtained by performing an extensive grid search over the complete set of strategies. Fig. 17 shows the performance of the JB algorithm against the two users that are randomly changing their power levels to overcome interference (this captures a much more difficult scenario as compared to standard adaptive mechanisms, such as power control schemes, in which the victim increases its power level until it reaches a maximum so as to overcome interference). Although each victim has a different adaption cycle (specifically, victim 1 changes its power levels based on the performance history over the past 50000 time instants and victim 2 adapts its power levels over a window of size 30000 time instants), the jammer is capable of tracking these changes in a satisfactory manner.

By using a weighted PER metric rather than a mean PER metric, the jammer can prioritize jamming one set of transmit-receive pairs against the others. Several other MU cases can easily be considered by using this framework. For example, by allowing the jammer to choose the direction of jamming as another action, the jammer can prioritize jamming only the transmit-receive pairs in a given direction rather than spread all its power uniformly across all directions. However, such improved jamming techniques

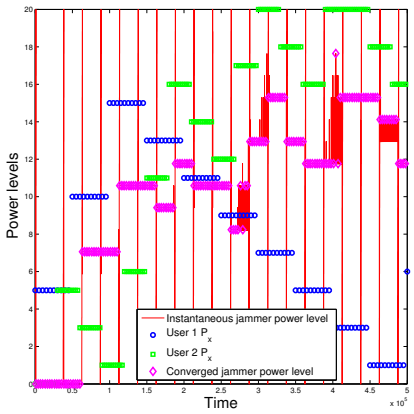

Fig. 17. PER achieved by the jammer against 2 stochastic users in the network. Both the users employ BPSK signaling scheme. The jammer learns to use the BPSK signaling scheme to achieve power efficient jamming strategies and also tracks the changes in the users' strategies.

will only come at the expense of more knowledge about the location of the users, users' behavior etc. Nevertheless, it is worth appreciating the applicability of the proposed algorithms to a wide variety of electronic warfare-type scenarios.

VI. Conclusion

In this paper, we proved that a cognitive jammer can learn the optimal physical layer jamming strategy in an electronic warfare-type scenario without having any *a priori* knowledge about the system dynamics. Novel learning algorithms based on the multi-armed bandit framework were developed to optimally jam victim transmitter-receiver pairs. The learning algorithms are capable of learning the optimal jamming strategies in both coherent and non-coherent scenarios where the jamming signal and the victim signal are either phase synchronous or asynchronous with each other. Also, the rate of learning is faster in comparison to commonly used reinforcement learning algorithms. These algorithms are capable of tracking the different strategies used by multiple adaptive transmitter-receiver pairs. Moreover, they come with strong theoretical guarantees on the performance including confidence bounds which are used to estimate the probability of successful jamming at any time instant.

References

[1] A. D. Wyner, "The wire-tap channel," *Bell System Technical Journal*, vol. 54, no. 8, pp. 1335-1387, Jan. 1975.

[2] I. Csiszar and J. Korner, "Broadcast channels with confidential messages," *IEEE Trans. Inf. Theory*, vol. 24, no. 3, pp. 339-348, May 1978.

[3] Y. Liang, H. V. Poor, and S. Shamai (Shitz), "Information theoretic security," *Foundations and Trends in Communications and Information Theory*, vol. 5, no. 7, pp. 355-580, 2008.

[4] T. Basar, "The Gaussian test channel with an intelligent jammer," *IEEE Trans. Inf. Theory*, vol. 29, no. 1, pp. 152-157, Jan. 1983.

[5] M. Medard, "Capacity of correlated jamming channels," In *Proc. Allerton Conf. Commun. Control and Comput.*, Monticello, IL, 1997, pp. 1043-1052.

[6] A. Kashyap, T. Basar, and R. Srikant, "Correlated jamming on MIMO Gaussian fading channels," *IEEE Trans. Inf. Theory*, vol. 50, no. 9, pp. 2119-2123, Sep. 2004.

[7] A. Mukherjee and A. L. Swindlehurst, "Jamming games in the MIMO wiretap channel with an active eavesdropper," *IEEE Trans. Signal Process*, vol. 62, no. 1, pp. 82-91, Jan. 2013.

[8] Y. O. Basciftci, C. E. Koksal and F. Ozguner, "To Obtain or not to Obtain CSI in the Presence of Hybrid Adversary", in *Proc. IEEE Intern. Symp. Inf. Theory*, Istanbul, Turkey, Jan. 2013, pp. 2865-2869.

[9] M. Azizoglu, "Convexity properties in binary detection problems," *IEEE Trans. Inf. Theory*, vol. 42, pp. 1316-1321, Jul. 1996.

[10] S. Bayram *et al.* "Optimum power allocation for average power constrained jammers in the presence of non-Gaussian noise," *IEEE Commun. Lett.,* vol. 16, no. 8, pp. 1153-1156, Aug. 2012.

[11] S. Amuru and R. M. Buehrer, "Optimal Jamming Strategies in Digital Communications-Impact of Modulation," in *Proc. Global Commun. Conf*, Austin, TX, Dec. 2014.

[12] K. Dabcevic, A. Betancourt, L. Marcenaro and C. S. Regazzoni, "A fictitious play-based game-theoretical approach to alleviating jamming attacks for cognitive radios", in *Proc. ICASSP*, Florence, Italy, May 2014.

[13] Y. E. Sagduyu, R. A. Berry and A. Ephremides, "Jamming games in wireless networks with incomplete information," *IEEE Commun. Mag.,* vol. 49, no. 8, pp. 112-118, Aug. 2011.

[14] R. McEliece and W. Stark, "An information theoretic study of communication in the presence of jamming", in *Proc. Int. Conf. Commun.*, 1981, pp. 45.3.1-45.3.5.

[15] S. Shamai (Shitz) and S. Verdú, "Worst-case power constrained noise for binary-input channels," *IEEE Trans. Inf. Theory,* vol. IT-38, no. 5, pp. 1494-1511, Sep. 1992.

[16] H. I. Volos and R. M. Buehrer, "Cognitive Engine Design for Link Adaptation: An Application to Multi-Antenna Systems," *IEEE Trans. Wireless Commun.,* vol. 9, no. 9, pp. 2902-2913, Sept. 2010.

[17] ——, "Cognitive Radio Engine Training," *IEEE Trans. Wireless Commun.,* vol. 11, no. 11, pp. 3878-3889, Nov. 2012.

[18] B. Wang, Y. Wu and K. J. R. Liu, "An anti-jamming stochastic game in cognitive radio networks," *IEEE J. Sel. Areas Commun.,* vol. 29, no. 4, pp. 877-889, Apr. 2011.

[19] Y. Wu, B. Wang , K. J. R. Liu and T. C. Clancy, "Anti-jamming games in multi-channel cognitive radio networks," *IEEE J. Sel. Areas Commun.,* vol. 30, no. 1, pp. 4-15, Jan. 2012.

[20] Y. L. Gwon, S. Dastangoo, C. E. Fossa, and H. T. Kung, "Competing Mobile Network Game: Embracing Antijamming and Jamming Strategies with Reinforcement Learning," in *Proc. Commun. Netw. Sec.,* Washington, D.C., Oct. 2013, pp. 28-36.

[21] C. Tekin and M. Liu, "Online learning in opportunistic spectrum access: a restless bandit approach," in *Proc. Intern. Conf. Comput. Commun.,*, Shanghai, China, Apr. 2011, pp. 2462-2470.

[22] S. Ahmad, M. Liu, T. Javidi, Q. Zhao and B. Krishnamachari, "Optimality of Myopic Sensing in Multichannel Opportunistic Access," *IEEE Trans. Inf. Theory,* vol. 55, pp. 4040-4050, Sept. 2009.

[23] Y. Gai, B. Krishnamachari and R. Jain, "Learning Multiuser Channel Allocations in Cognitive Radio Networks: A Combinatorial Multi-Armed Bandit Formulation," in *Proc. DYSPAN*, Singapore, April 2010.

[24] Q. Wang, P. Xu, K. Ren, and X.-Y. Li, "Towards optimal adaptive ufh-based anti-jamming wireless communication," *IEEE J. Sel. Areas Commun.*, vol. 30, no. 1, pp. 16-30, Jan. 2012.

[25] N. Gulati and K. R. Dandekar, "Learning State Selection for Reconfigurable Antennas: A Multi-Armed Bandit Approach", *IEEE Trans. Antenna Propag.,* vol. 62, no. 3, pp. 1027-1038, Mar. 2014.

[26] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite-time analysis of the multi-armed bandit problem," *Machine Learning*, vol. 47, no. 2, pp. 235-256, May 2002.

[27] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire, "The non-stochastic multi-armed bandit problem", *SIAM J. Comput.,* vol. 32, no. 1, pp. 48-77, Jan. 2002.

[28] R. Kleinberg, "Nearly tight bounds for the continuum-armed bandit problem," in *Proc. Neural Inf. Proc. Syst.,* 2004.

[29] 3GPP, TS 36.201, "Evolved Universal Terrestrial Radio Access (E-UTRA); LTE Physical Layer-General Description (Release 8)".

[30] S. Amuru and R. M. Buehrer, "Optimal Jamming using Delayed Learning," in *Proc. Military Commun. Conf.,* Baltimore, MD, Oct. 2014.

[31] Z. Serceki and M. Wilhoyte, "Method for determining packet error rate of wireless lan stations," U.S. Patent 20040076138A1, Apr. 22, 2004.

[32] J-Y. Audibert, R. Munos, and C. Szepesvari, "Exploration-exploitation trade-off using variance estimates in multi-armed bandits," *Theor. Comput. Sci.,* vol. 410, no. 19, pp. 1876-1902, Apr. 2009.

[33] R. A. Poisel, *Introduction to Communication Electronic Warfare Systems*. Artech House, 2008.

[34] P. Auer and R. Ortner, "UCB revisited: Improved regret bounds for the stochastic multi-armed bandit problem," *Periodica Mathematica* Hungarica, vol. 61, pp. 55-65, 2010.

[35] M. Strasser, S. Čapkun, and M. Čagalj, "Jamming-resistant key establishment using uncoordinated frequency hopping", in *Proc. IEEE Symp. Security and Privacy*, Oakland, CA, May 2008, pp. 64-78.

[36] C. Tekin , L. Canzian, and M. van der Schaar, "Context adaptive Big Data stream mining", in *Proc. Allerton Conf. Commun. Control and Comput.,* Monticello, IL, Oct. 2014.

APPENDIX A
PROOF OF THEOREM 1

For the system model in Section II, the average probability of error at the victim receiver that uses a maximum likelihood (ML) detector (since it is assumed that the victim transmit-receive pair is not aware of the presence of the jammer) is given by

$$p_e\left(j, \mathrm{SNR}, \mathrm{JNR}\right) = 1 - \int_x \int_\Omega f_N\left(y - \sqrt{\mathrm{SNR}}x - \sqrt{\mathrm{JNR}}j\right) f_X(x) dy dx, \tag{5}$$

where Ω indicates the ML decision region for x. For instance, when the signal levels are $\pm A$, $\Omega = \mathsf{real}\left(y\right) < 0$ when $x = -A$ and $\Omega = \mathsf{real}\left(y\right) > 0$ when $x = +A$. In the above equation, the received signal normalized by the noise power σ^2 is considered. Further, f_X indicates the distribution of the signal x (described by the modulation scheme used by the victim) and f_N indicates the additive white Gaussian noise distribution. For a pulsed jamming signal with pulsing ratio ρ, the SER is given by $\rho p_e(j, \mathrm{SNR}, \frac{\mathrm{JNR}}{\rho}) + (1-\rho)p_e(j, \mathrm{SNR}, 0)$. We first establish the Hölder continuity of $p_e\left(j, \mathrm{SNR}, \mathrm{JNR}\right)$ which can then be used to prove the Hölder continuity of p_e for pulsed jamming scenarios.

In order to prove that SER i.e., p_e is uniformly locally Lipschitz, we show that $|p_e(j, \mathrm{SNR}, \mathrm{JNR}_1) - p_e(j, \mathrm{SNR}, \mathrm{JNR}_2)| \leq L|\mathrm{JNR}_1 - \mathrm{JNR}_2|^\alpha$ for some $L > 0$ and $\alpha > 0$. Using (5) we have,

$$\begin{aligned}
&|p_e(j, \mathrm{SNR}, \mathrm{JNR}_1) - p_e(j, \mathrm{SNR}, \mathrm{JNR}_2)| \\
&= |\int_x \int_\Omega \left[f_N\left(y - \sqrt{\mathrm{SNR}}x - \sqrt{\mathrm{JNR}_2}j\right) - f_N\left(y - \sqrt{\mathrm{SNR}}x - \sqrt{\mathrm{JNR}_1}j\right)\right] f_X(x) dy dx| \\
&\leq \int_x \int_\Omega \left[|f_N\left(y - \sqrt{\mathrm{SNR}}x - \sqrt{\mathrm{JNR}_2}j\right) - f_N\left(y - \sqrt{\mathrm{SNR}}x - \sqrt{\mathrm{JNR}_1}j\right)|\right] f_X(x) dy dx.
\end{aligned} \tag{6}$$

Thus it is sufficient to show that $|f_N\left(y - \sqrt{\mathrm{SNR}}x - \sqrt{\mathrm{JNR}_2}j\right) - f_N\left(y - \sqrt{\mathrm{SNR}}x - \sqrt{\mathrm{JNR}_1}j\right)| \leq L'|\mathrm{JNR}_1 - \mathrm{JNR}_2|^{\alpha'}$ for some $L' > 0$ and $\alpha' > 0$ which follows from the definition of f_N as it is the probability density function (pdf) of the noise signal n. We briefly show it below for completeness. Since we already normalized the signal by σ^2, the pdf of n is now given by the zero mean unit variance

Gaussian distribution.

$$
\begin{aligned}
&|f_N\left(y-\sqrt{\mathrm{SNR}}x-\sqrt{\mathrm{JNR}_2}j\right)-f_N\left(y-\sqrt{\mathrm{SNR}}s-\sqrt{\mathrm{JNR}_1}j\right)|\\
&=|\frac{1}{\sqrt{2\pi}}\left[\exp(-(y-\sqrt{\mathrm{SNR}}x-\sqrt{\mathrm{JNR}_2}j)^2)-\exp(-(y-\sqrt{\mathrm{SNR}}x-\sqrt{\mathrm{JNR}_1}j)^2)\right]|\\
&\approx\frac{1}{\sqrt{2\pi}}\left[|(\sqrt{JNR}_1-\sqrt{JNR}_2)j|\right], \qquad (7)
\end{aligned}
$$

where the last approximation is obtained by ignoring the higher order terms since we only consider the cases where $|\mathrm{JNR}_1-\mathrm{JNR}_2|\leq\delta$ i.e., local Hölder continuity. Then, for a given jamming signal j, (7) can be bounded as

$$
\begin{aligned}
\frac{1}{\sqrt{2\pi}}\left[|(\sqrt{JNR}_1-\sqrt{JNR}_2)j|\right]&\leq\sqrt{\frac{\mathrm{JNR}_2}{2\pi}}|\left(\sqrt{1+\frac{\delta}{\mathrm{JNR}_2}}-1\right)|\approx\sqrt{\frac{\mathrm{JNR}_2}{2\pi}}|\left(1+\frac{\delta}{2\mathrm{JNR}_2}-1\right)|\\
&\leq\sqrt{\frac{1}{2\pi\mathrm{JNR}_{\min}}}\delta=\sqrt{\frac{1}{2\pi\mathrm{JNR}_{\min}}}(\mathrm{JNR}_1-\mathrm{JNR}_2), \qquad (8)
\end{aligned}
$$

which proves that argument inside the integral in (6) is uniformly locally Lipschitz with $L'=\sqrt{\frac{1}{2\pi\mathrm{JNR}_{\min}}}$ and $\alpha'=1$. In the above proof we used the fact that $|j|\leq 1$ for standard signaling schemes that are employed by the jammer (for the AWGN jamming signal, the SER is obtained by using a Gaussian distribution with variance $1+\mathrm{JNR}$ i.e., a slightly different approach when compared to (5) is taken and by following the above sequence of arguments, Hölder continuity can be proved even in this case). Using (8), the overall SER i.e., $p_e(j,\mathrm{SNR},\mathrm{JNR})$ is also uniformly locally Hölder continuous. By following the same steps, the Hölder continuity for the pulsed jamming cases can also be proved. An example for the Hölder continuity in the pulsed jamming case is shown in Section III.

APPENDIX B
PROOF OF THEOREM 2

Since the set of signaling schemes is discrete, we first obtain the regret bound for a particular signaling scheme $\mathcal{J}$. It is easy to see that the overall regret bound is a scaled version (by N_{mod}) of the regret achievable for a single signaling scheme. Since the time horizon of the inner loop of Algorithm 1 is T, we first show that the regret incurred by the inner loop is $\mathcal{O}(\sqrt{M^2T\log(T)})$. Since the overall time horizon is generally unknown, the algorithm is run for several rounds of time steps on the order of 2^i as shown in Algorithm 1 and the regret bounds for the overall algorithm can be achieved by using the doubling trick [32].

The upper bound on the overall regret incurred by Algorithm 1 can be obtained by upper bounding $\sum_{t=1}^{T}\left(\bar{C}(\mathcal{J},\mathbf{s}^*)-\bar{C}(\mathcal{J},\mathbf{s}_t)\right)$, where $\bar{C}$ indicates the average cost function and $\mathbf{s}^*$ is the best strategy for a given signaling scheme $\mathcal{J}$ and $\mathbf{s}_t$ is the actual strategy chosen at time t. For ease of presentation, $\mathcal{J}$

is ignored in the rest of the proof. We obtain the regret bound in two steps by rewriting it as

$$\sum_{t=1}^{T}\left(\bar{C}(\mathbf{s}^*)-\bar{C}(\mathbf{s}_t)\right)=\sum_{t=1}^{T}\left(\bar{C}(\mathbf{s}^*)-\bar{C}(\mathbf{s}')\right)+\sum_{t=1}^{T}\left(\bar{C}(\mathbf{s}')-\bar{C}(\mathbf{s}_t)\right), \tag{9}$$

where $\mathbf{s}' \in \{1/M, 2/M, \ldots, 1\}\times \mathrm{JNR}_{\min}+(\mathrm{JNR}_{\max}-\mathrm{JNR}_{\min})*\{1/M, 2/M, \ldots, 1\}$ is the strategy nearest (in terms of the Euclidean distance) to $\mathbf{s}^*$. Then we have $||\mathbf{s}'-\mathbf{s}^*||=\sqrt{(\mathrm{JNR}'-\mathrm{JNR}^*)^2+(\rho'-\rho^*)^2}$ $\leq \sqrt{\frac{2}{M^2}}$ based on the discretization of the continuous arms set in Algorithm 1.

For the first term in the above equation, by using the Hölder continuity properties of the average cost function $\bar{C}$, we have

$$\mathbf{E}\Big(\sum_{t=1}^{T}C_t(\mathbf{s}^*)-C_t(\mathbf{s}')\Big)=\sum_{t=1}^{T}\left(\bar{C}(\mathbf{s}^*)-\bar{C}(\mathbf{s}')\right)\leq TL\Big(\frac{2}{M^2}\Big)^{\alpha/2}. \tag{10}$$

We now bound the second term $\mathbf{E}\Big(\sum_{t=1}^{T}C_t(\mathbf{s}')-C_t(\mathbf{s}_t)\Big)=\sum_{t=1}^{T}\left(\bar{C}(\mathbf{s}')-\bar{C}(\mathbf{s}_t)\right)$. Due to the discretization technique used in Algorithm 1, this problem is equivalent to a standard MAB problem with M^2 arms [26]. In order to bound (10), we define two sets of arms: near-optimal arms and sub-optimal arms. We set $\Delta=\sqrt{M^2\log(T)/T}$ and say that an arm is sub-optimal in this case, if its regret incurred is greater than Δ and near-optimal when its regret is less than Δ. Thus, for a near-optimal arm, even when that arm is selected at all time steps, the contribution to regret will be at most $T\Delta$. In contrast for a sub-optimal arm, the contribution to the regret when it is selected can be large. Since we use the UCB1 algorithm, it can be shown that the sub-optimal arms will be chosen only $\mathcal{O}(\log(T)/\Delta(\mathbf{s})^2)$ times ($\Delta(\mathbf{s})$ is the regret of the strategy $\mathbf{s}$) [26], before they are identified as sub-optimal. Thus the regret for these sub-optimal arms is on the order of $\mathcal{O}(\log(T)/\Delta)$ since $\Delta(\mathbf{s})>\Delta$. From these arguments the second term in (9) can be upper bounded as

$$\mathbf{E}\Big(\sum_{t=1}^{T}C_t(\mathbf{s}_t)-C_t(\mathbf{s}')\Big)\leq\mathcal{O}(\sqrt{M^2T\log(T)}). \tag{11}$$

Using (10) and (11), and setting $M=\lceil(\sqrt{\frac{T}{\log(T)}}L2^{\alpha/2})^{\frac{1}{1+\alpha}}\rceil$ (this is obtained by matching the regret bounds shown in (10) and (11)), the regret for any given signaling scheme is given by $\mathcal{O}(\sqrt{M^2T\log(T)})$. By noting the fact that the jammer can choose from N_{mod} possible signaling schemes and using the value of M, the doubling trick, and summing the regret over all inner loop iterations of Algorithm 1, the regret over the entire time horizon n can be expressed as $\mathcal{O}(N_{mod}n^{\frac{\alpha+2}{2(\alpha+1)}}(\log n)^{\frac{\alpha}{2(\alpha+1)}})$.

APPENDIX C

HIGH CONFIDENCE BOUNDS

Here, we present high confidence bounds on the one-step and cumulative regret.

A. *Proof of Theorem 3*

We bound the one-step regret as follows,

$$\begin{aligned}
P(\bar{C}(\mathcal{J}^*, \mathbf{s}^*) - \bar{C}(\mathcal{J}_t, \mathbf{s}_t) > \delta) &= P(\bar{C}(\mathcal{J}^*, \mathbf{s}^*) - \bar{C}(\mathcal{J}_t, \mathbf{s}^*) + \bar{C}(\mathcal{J}_t, \mathbf{s}^*) - \bar{C}(\mathcal{J}_t, \mathbf{s}_t) > \delta) \\
&= P(\bar{C}(\mathcal{J}^*, \mathbf{s}^*) - \bar{C}(\mathcal{J}_t, \mathbf{s}^*) + \bar{C}(\mathcal{J}_t, \mathbf{s}^*) - \bar{C}(\mathcal{J}_t, \mathbf{s}') + \bar{C}(\mathcal{J}_t, \mathbf{s}') - \bar{C}(\mathcal{J}_t, \mathbf{s}_t) > \delta) \\
&\leq P(\bar{C}(\mathcal{J}^*, \mathbf{s}^*) - \bar{C}(\mathcal{J}_t, \mathbf{s}^*) \geq \frac{\delta}{4}) + P(\bar{C}(\mathcal{J}_t, \mathbf{s}^*) - \bar{C}(\mathcal{J}_t, \mathbf{s}') \geq \frac{\delta}{2}) \\
&\qquad + P(\bar{C}(\mathcal{J}_t, \mathbf{s}') - \bar{C}(\mathcal{J}_t, \mathbf{s}_t) \geq \frac{\delta}{4}) \qquad (12)
\end{aligned}$$

For the second term in (12), we have

$$P(\bar{C}(\mathcal{J}_t, \mathbf{s}^*) - \bar{C}(\mathcal{J}_t, \mathbf{s}') \geq \frac{\delta}{2}) = 1 - P(\bar{C}(\mathcal{J}_t, \mathbf{s}^*) - \bar{C}(\mathcal{J}_t, \mathbf{s}') \leq \frac{\delta}{2}), \qquad (13)$$

which for $\delta > 2L(\frac{2}{M^2})^{\frac{\alpha}{2}}$ converges in probability to 0 because

$$1 - P(\bar{C}(\mathcal{J}_t, \mathbf{s}^*) - \bar{C}(\mathcal{J}_t, \mathbf{s}') \leq \frac{\delta}{2}) \leq 1 - P(\bar{C}(\mathcal{J}_t, \mathbf{s}^*) - \bar{C}(\mathcal{J}_t, \mathbf{s}') \leq L|\mathbf{s}_t - \mathbf{s}'|^{\alpha} \leq \frac{\delta}{2}). \qquad (14)$$

The last equality is a result of the Hölder continuity properties of the cost function.

Recall that we use the UCB1 algorithm to choose arms within one round of JB. Hence to bound the term $P(\bar{C}(\mathcal{J}^*, \mathbf{s}^*) - \bar{C}(\mathcal{J}_t, \mathbf{s}^*) \geq \frac{\delta}{4})$ in (12), we define two sets of arms (i) set of arms $\mathcal{J}_>$ with sub-optimality gap $\Delta_i^{\mathcal{J}} = \bar{C}(\mathcal{J}^*, \mathbf{s}^*) - \bar{C}(\mathcal{J}^i, \mathbf{s}^*) > \delta/4$ for all $1 \leq i \leq N_{mod}$, also referred to as sub-optimal arms and (ii) rest of the arms denoted by $\mathcal{J}_<$. Then we have

$$P(\bar{C}(\mathcal{J}^*, \mathbf{s}^*) - \bar{C}(\mathcal{J}_t, \mathbf{s}^*) \geq \frac{\delta}{4}) = P(\cup_{i \in \boldsymbol{\mathcal{J}}_>} \mathsf{Arm}_t = \mathcal{J}^i) \leq \sum_{i \in \boldsymbol{\mathcal{J}}_>} P(\mathsf{Arm}_t = \mathcal{J}^i), \qquad (15)$$

where $\mathsf{Arm}_t = \mathcal{J}^i$ indicates that $\mathcal{J}^i$ is chosen at time t and the inequality follows from the union bound. Let $J(T)$ be the (random) set of time steps for which $j_i(T) \leq 8 \log T / (\Delta_i^{\mathcal{J}})^2$ for some sub-optimal signaling scheme i in $\{1, 2, \ldots, T\}$. The time steps in $J(T)$ are the time steps in which there is a high probability that a sub-optimal signaling scheme is selected. In contrast, for time steps in $J(T)^c := \{1, 2, \ldots, T\} \backslash J(T)$, all the arms are selected sufficiently many times to have accurate estimates so that the best arm can be correctly identified with a very high probability. Following the analysis of Theorem 1 in [26] for the UCB1 algorithm, it can be shown that at time steps $t \in J(T)^c$, the probability

of choosing a sub-optimal signaling scheme i is bounded above by $2t^{-4}$. Hence, for $t \in J(T)^c$, (16) can be written as

$$P(\bar{C}(\mathcal{J}^*, \mathbf{s}^*) - \bar{C}(\mathcal{J}_t, \mathbf{s}^*) \geq \frac{\delta}{4}) \leq \sum_{i \in \boldsymbol{\mathcal{J}}_>} 2t^{-4} < 2N_{mod}t^{-4}, \tag{16}$$

Along similar lines, for the third term in (12), we have

$$P(\bar{C}(\mathcal{J}_t, \mathbf{s}') - \bar{C}(\mathcal{J}_t, \mathbf{s}_t) \geq \frac{\delta}{4}) < 2M^2t^{-4}, \tag{17}$$

for all $t \in [1, T]\backslash S(T)$ where $\mathcal{S}_>$ and $S(T)$ are defined along similar lines as $\mathcal{J}_>$ and $J(T)$ but for the M^2 discrete arms that correspond to the discretized space of JNR and ρ.

Overall, by choosing $\delta > 2L(\frac{2}{M^2})^{\frac{\alpha}{2}}$ (specifically, we choose $\delta = 4L(\frac{2}{M^2})^{\frac{\alpha}{2}}$), the one step regret can be shown to converge in probability to 0 as $P\left(\bar{C}(\mathcal{J}^*, \mathbf{s}^*) - \bar{C}(\mathcal{J}_t, \mathbf{s}_t) > \delta\right) \leq 2(N_{mod} + M^2)t^{-4}\ \forall t \in [1, T]\backslash\{\mathcal{J}(T) \cup \mathcal{S}(T)\}$. Since the one-step regret converges in probability to 0, it implies that the jammer can reach the optimal value of the cost function by using Algorithm 1. The value of δ is governed by the discretization M and the regret incurred by the jammer is mainly due to discretization of the continuum arm space which for larger values of M (possible as T increases as shown in Algorithm 1) tightens the confidence bound on the regret.

While the above analysis presented a confidence bound on the one-step regret, it did so by splitting the choice of the signaling scheme and the continuous parameters. We unify these analyses below. The overall sub-optimality gap of the ith arm $1 \leq i \leq N_{mod}M^2$, denoted by $\{\mathcal{J}^i, \mathbf{s}^i\}$ (recall that $N_{mod}M^2$ arms can be chosen in one round of JB), is defined as $\bar{C}(\mathcal{J}^*, \mathbf{s}^*) - \bar{C}(\mathcal{J}^i, \mathbf{s}^i)$. Let $u_i(t)$ denote the total number of times the ith arm has been chosen until time t and $U(T)$ indicate the set of time instants $t \in [1, T]$ for which $u_i(t) \leq \frac{8\log(T)}{\Delta_i^2}$ for some sub-optimal arm $i \in \mathcal{U}_>$ with a sub-optimality gap Δ_i. It is easy to see that these definitions directly result in the following relationships: $U_> \subseteq \mathcal{J}_> \cup \mathcal{S}_>$ and $U(T) \subseteq J(T) \cup S(T)$.

Note that if the jammer knows a lower bound $\underline{\Delta} > 0$ on the minimum sub-optimality gap (for instance, in a wireless communication setting when SER is taken as the cost function, the lower bound can indicate the smallest error in the SER that JB can be allowed to make i.e., we want the choices of JB to be closest to the optimal SER^* that can be achievable) i.e., $\Delta_{\min} := \min_{i \in \mathcal{U}_>} \Delta_i$, then it can estimate if t belongs to $U(T)$ or not by checking if $u_i(t) > 8\log T/\underline{\Delta}^2$ for all $i = 1, 2, \ldots, N_{mod}$. Let $\hat{U}(T)^c$ be the set of this estimated time slots, for which we have $\hat{U}(T)^c \subseteq U(T)^c$. The jammer will know that for at least $t \in \hat{U}(T)^c$ the proposed confidence bound will hold true. While the exact set of all time

instants that belong to $U(T)$ may be unknown to the jammer, we can bound the size of the set $E[U(T)]$ as follows [26]:

$$E[|U(T)|] \leq \sum_{t=1}^{T} P(\text{a sub-optimal arm } i \in \mathcal{U}_{>} \text{ is chosen at } t)$$
$$\leq 8 \sum_{i \in \mathcal{U}_{>}} \left(\frac{\log T}{\Delta_i^2} \right) + \left(1 + \frac{\pi^2}{3} \right) |U_{>}|,$$

which follows from Theorem 1 in [26]. This suggests that our confidence bounds hold in all except logarithmically many time slots in expectation. We stress here that the size of the set $U(T)$ presented above is the worst case bound and the algorithm performs much better than predicted by the bounds as shown in Fig. 10 in Section V. This completes the proof of the Theorem.

B. Proof of Theorem 4

Here, we present a high confidence bound on the cumulative regret incurred by the jammer when it uses Algorithm 1. Similar to the regret bound in Theorem 2, we have for any given signaling scheme used by the jammer,

$$P\left(\sum_{t=1}^{T} \left(\bar{C}(\mathbf{s}^*) - \bar{C}(\mathbf{s}_t) \right) > \delta \right) = P\left(\sum_{t=1}^{T} \left(\bar{C}(\mathbf{s}^*) - \bar{C}(\mathbf{s}') + \bar{C}(\mathbf{s}') - \bar{C}(\mathbf{s}_t) \right) > \delta \right)$$
$$\leq P\left(\sum_{t=1}^{T} \left(\bar{C}(\mathbf{s}^*) - \bar{C}(\mathbf{s}') \right) > \delta/2 \right) + P\left(\sum_{t=1}^{T} \left(\bar{C}(\mathbf{s}') - \bar{C}(\mathbf{s}_t) \right) > \delta/2 \right), \tag{18}$$

where $\mathbf{s}'$ was defined earlier. For the first term, choosing $\delta > 2TL \left(\frac{2}{M^2} \right)^{\alpha/2} = O(T^{\frac{1+2\alpha}{2(1+\alpha)}} (\log T)^{\frac{1}{2(1+\alpha)}})$ and using the Hölder continuity properties of the cost function gives the following,

$$P\left(\sum_{t=1}^{T} \left(\bar{C}(\mathbf{s}^*) - \bar{C}(\mathbf{s}') \right) > \delta/2 \right) = 1 - \left[P\left(\sum_{t=1}^{T} \left(\bar{C}(\mathbf{s}^*) - \bar{C}(\mathbf{s}') \right) < \delta/2 \right) \right]$$
$$\leq 1 - \left[P\left(\sum_{t=1}^{T} \left(\bar{C}(\mathbf{s}^*) - \bar{C}(\mathbf{s}') \right) < LT|\mathbf{s}^* - \mathbf{s}'|^{\alpha} < \delta/2 \right) \right] = 0, \ \forall \epsilon > 0, \tag{19}$$

where the last equality is a result of the choice of δ.

For the second term in (18), as earlier, we use tricks from the UCB1 algorithm to obtain confidence bounds as it is a standard finite-armed (M^2 arms) bandit problem [32]. Since the regret in the case of a finite-armed bandit problem can be represented in terms of the number of times an arm has been chosen, we first present a lemma that provides confidence bounds on the same i.e., we evaluate $P(s_k(T) > u)$ where $s_k(T)$ denotes the total number of times arm $k (1 \leq k \leq M^2)$ has been chosen until time T.

Consider the event $\mathcal{E}$ defined by

$$\mathcal{E} = \{\{\forall t : u+1 \leq t \leq T \text{ s.t. } B_{k,u,t} \leq \tau\} \ \wedge \ \{\forall v : 1 \leq v \leq T-u \text{ s.t. } B_{k^*,v,u+v} > \tau\}\}, \tag{20}$$

where $B_{k,u,t}$ indicates the UCB1 index [26] of arm k at time instant t when it has been pulled u times until the time time instant t, $B_{k^*,v,u+v}$ indicates the UCB1 index of the optimal arm defined along similar lines and $\tau \in \mathbb{R}$ is any real number. In the definition of event $\mathcal{E}$, we have two sub-events, where the first one suggests that the UCB1 index of the kth arm is less than τ and the second event indicates that since arm k was chosen u times, the optimal arm will be chosen $v \in [1, T-u]$ times when its index exceeds τ. Then, for any $1 \leq v \leq T-u$ and $u+v \leq t \leq T$ we have $B_{k^*,v,t} \geq B_{k^*,v,u+v} > \tau \geq B_{k,u,t}$ (because we want to compare the indices of the arms at time t given the fact that arm k was chosen u times before this and the optimal arm was chosen v times before this time instant t), which suggests that arm k would not be chosen the $(u+1)$th time at any time $t \leq T$ (notice that the exact time instants at which the arms were chosen do not matter, only the number of times an arm chosen decides the UCB1 index of that arm). Thus we have by contradiction

$$\begin{aligned} P(s_k(T) > u) \leq P\left(\exists t : \ u+1 \leq t \leq T \text{ s.t. } B_{k,u,t} > \tau\right) \\ + P\left(\exists v : \ 1 \leq v \leq T-u \text{ s.t. } B_{k^*,v,u+v} \leq \tau\right). \end{aligned} \tag{21}$$

We will now bound the second term in (18) i.e., $P\left(\sum_{t=1}^{T}\left(\bar{C}(\mathbf{s}') - \bar{C}(\mathbf{s}_t)\right) > \delta/2\right)$. As mentioned earlier, since this is a standard MAB problem, we have the following

$$\begin{aligned} P\left(\sum_{t=1}^{T}\left(\bar{C}(\mathbf{s}') - \bar{C}(\mathbf{s}_t)\right) > \delta/2\right) &= P\left(\sum_{k:\Delta_k^{\mathcal{S}}>0} \Delta_k^{\mathcal{S}} s_k(T) > \delta/2\right) \leq \sum_{k:\Delta_k^{\mathcal{S}}>0} P\left(s_k(T) > \frac{\delta}{2\Delta_k^{\mathcal{S}}}\right) \\ &\leq \sum_{k:\Delta_k^{\mathcal{S}}>0} P\left(s_k(T) > \frac{\delta}{2M^2\Delta_k^{\mathcal{S}}}\right), \end{aligned} \tag{22}$$

where $\Delta_k^{\mathcal{S}} = \bar{C}(\mathbf{s}') - \bar{C}(\mathbf{s}^k)$ i.e., the regret incurred for playing arm $k \in [1, M^2]$. For any $\delta \geq 16M^2 \frac{\log(T)}{\Delta_{\min}^{\mathcal{S}}}$ where $\Delta_{\min}^{\mathcal{S}}$ is the minimum regret incurred across all arms, we have $\frac{\delta}{2M^2\Delta_k^{\mathcal{S}}} \geq \frac{8}{(\Delta_k^{\mathcal{S}})^2}\log(T)$. Let $u_k = \lceil \frac{\delta}{2M^2\Delta_k^{\mathcal{S}}} \rceil$. Using the bound on $s_k(T)$ in (21) and with $u = u_k$, $\tau = \bar{\mathbf{s}}'$, we have

$$P(s_k(T) > u_k) \leq \sum_{t=u_k+1}^{T} P(B_{k,u_k,t} > \bar{C}(\mathbf{s}')) + \sum_{v=1}^{T-u_k} P(B_{k^*,v,u_k+v} \leq \bar{C}(\mathbf{s}')). \tag{23}$$

Using the condition on δ, we have $\sqrt{\frac{2\log(T)}{u_k}} \leq \Delta_k^{\mathcal{S}}/2$ which upon rearranging the terms also gives $T \leq e^{\frac{u_k(\Delta_k^{\mathcal{S}})^2}{8}}$. Then we can bound $P(B_{k,u_k,t} > \bar{C}(\mathbf{s}'))$ (which indicates the fact that the upper confidence

bound on the sub-optimal arm is higher than the mean reward/cost function incurred by the optimal strategy $\mathbf{s}'$) as

$$P(B_{k,u_k,t} > \bar{C}(\mathbf{s}')) \leq P(B_{k,u_k,T} > \bar{C}(\mathbf{s}')) = P\left(\hat{C}(\mathbf{s}^k) + \sqrt{\frac{2\log(T)}{u_k}} > \bar{C}(\mathbf{s}^k) + \Delta_k^{\mathcal{S}}\right)$$
$$\leq P\left(\bar{C}(\mathbf{s}^k) - \bar{C}(\mathbf{s}') > \Delta_k^{\mathcal{S}}/2\right) \leq e^{-u_k(\Delta_k^{\mathcal{S}})^2/2}, \tag{24}$$

where $\hat{C}(\mathbf{s}^k)$ indicates the estimate of the cost function/reward obtained by suing strategy $\mathbf{s}^k$, the second inequality is obtained by using the definition of u_k and δ, and the third inequality by using the Chernoff-Hoeffding bound. Using (24) and the fact that $T \leq e^{\frac{u_k(\Delta_k^{\mathcal{S}})^2}{8}}$, the first summation in (23) can be upper bounded by $e^{-\frac{u_k(\Delta_k^{\mathcal{S}})^2}{4}}$.

The second summation in (23) is bounded by using the Chernoff-Hoeffding bound as $\sum_{v=1}^{T-u_k} P(B_{k^*,v,u_k+v} \leq \bar{C}(\mathbf{s}')) \leq \sum_{v=1}^{T-u_k}(u_k+v)^{-4} \leq \sum_{y=u_k}^{T} y^{-4} \leq \int_{y=u_k}^{\infty} y^{-4} = \frac{u_k^{-3}}{3}$ (by change of variable). Overall, (23) can be upper bounded as

$$P(s_k(T) > u_k) \leq e^{-\frac{u_k(\Delta_k^{\mathcal{S}})^2}{4}} + \frac{u_k^{-3}}{3}. \tag{25}$$

For $u_k = \lceil \frac{\delta}{2M^2\Delta_k^{\mathcal{S}}} \rceil$, the above bound is given by

$$P\left(s_k(T) > \frac{\delta}{2M^2\Delta_k^{\mathcal{S}}}\right) = P\left(s_k(T) > \lceil \frac{\delta}{2M^2\Delta_k^{\mathcal{S}}} \rceil\right) \leq e^{-\frac{\delta\Delta_k^{\mathcal{S}}}{8M^2}} + 8\delta^{-3}\frac{(M^2\Delta_k^{\mathcal{S}})^3}{3}. \tag{26}$$

Thus we have the upper bound on the second term in (18) as

$$P\left(\sum_{t=1}^{T}\left(\bar{C}(\mathbf{s}') - \bar{C}(\mathbf{s}_t)\right) > \delta/2\right) \leq \sum_{k:\Delta_k>0}\left\{e^{-\frac{\delta\Delta_k}{8M^2}} + 8\delta^{-3}\frac{(M^2\Delta_k^{\mathcal{S}})^3}{3}\right\}$$
$$\leq (M^2-1)\left(1 + \frac{8}{3}\delta^{-3}M^6\right) \approx \frac{8}{3}\delta^{-3}M^8 \triangleq \epsilon, \tag{27}$$

where we used the fact that $\Delta_k \in [0,1]$. Overall using (19) and (27), for all $\delta > \max\left(2TL\left(\frac{2}{M^2}\right)^{\alpha/2}, 16M^2\frac{\log(T)}{\Delta_{\min}^{\mathcal{S}}}\right)$, we have

$$P\Big(\sum_{t=1}^{T}\Big(\bar{C}(\mathbf{s}^*) - \bar{C}(\mathbf{s}_t)\Big) > \left(\frac{8}{3\epsilon}\Big(\frac{T}{\log(T)}\Big)^{\frac{4}{1+\alpha}}\right)^{1/3}\Big) < \epsilon. \tag{28}$$

Since $\Delta_{\min}$ is unknown *a priori*, a jammer can use any known lower bound $\underline{\Delta}$ to obtain δ. This lower bound can be obtained as described in the proof of Theorem 3.

Similar bounds i.e., of the same order can be obtained irrespective of the signaling scheme used by the jammer (here the jammer uses N_{mod} signaling schemes) as it is just another finite armed bandit problem.

This can be done as follows

$$P\Big(\sum_{t=1}^{T}(\bar{C}(\mathcal{J}^*,\mathbf{s}^*)-\bar{C}(\mathcal{J}_t,\mathbf{s}_t)>\delta)\Big)=P\Big(\sum_{t=1}^{T}(\bar{C}(\mathcal{J}^*,\mathbf{s}^*)-\bar{C}(\mathcal{J}^*,\mathbf{s}_t)+\bar{C}(\mathcal{J}^*,\mathbf{s}_t)-\bar{C}(\mathcal{J}_t,\mathbf{s}_t)>\delta)\Big)$$
$$\leq P\Big(\sum_{t=1}^{T}(\bar{C}(\mathcal{J}^*,\mathbf{s}^*)-\bar{C}(\mathcal{J}^*,\mathbf{s}_t)>\delta/2)\Big)+P\Big(\sum_{t=1}^{T}(\bar{C}(\mathcal{J}^*,\mathbf{s}_t)-\bar{C}(\mathcal{J}_t,\mathbf{s}_t)>\delta/2)\Big), \quad (29)$$

where the first term in the above equation can be directly bounded by using the analysis presented for (28) and the second term can be bounded directly using the properties of the UCB1 algorithm since it is a discrete arm selection (specifically, by using the steps in (20)-(28)).

APPENDIX D
HIGH CONFIDENCE BOUND ON ESTIMATES

*P*roof of Theorem 5: A high confidence bound on the mean estimate of the reward/cost function for any strategy that is used at time t by the jammer is developed. To do so, we evaluate $P(\bar{C}(\mathcal{J}^*,\mathbf{s}^*)-\hat{C}(\mathcal{J}_t,\mathbf{s}_t)>\delta)$ as follows,

$$P(\bar{C}(\mathcal{J}^*,\mathbf{s}^*)-\hat{C}(\mathcal{J}_t,\mathbf{s}_t)>\delta)\leq P(\bar{C}(\mathcal{J}^*,\mathbf{s}^*)-\bar{C}(\mathcal{J}_t,\mathbf{s}_t)>\frac{\delta}{2})+P(\bar{C}(\mathcal{J}_t,\mathbf{s}_t)-\hat{C}(\mathcal{J}_t,\mathbf{s}_t)>\frac{\delta}{2}), \quad (30)$$

where $\bar{C}(\mathcal{J}_t,\mathbf{s}_t)$ is the actual mean reward/cost of the strategy $(\mathcal{J}_t,\mathbf{s}_t)$. The first term can be bounded using Theorem 3 where it can be shown to be less than $2(N_{mod}+M^2)t^{-4}$ for all $\delta>2^{\frac{5\alpha+4}{2(1+\alpha)}}L^{\frac{1}{1+\alpha}}\left(\frac{\log T}{T}\right)^{\frac{\alpha}{2(1+\alpha)}}$ and $t\in[1,T]\backslash\{U(T)\}$, where $U(T)$ is defined in the proof of Theorem 3. For the second term, notice that we are comparing the actual and estimated mean rewards of the strategy $(\mathcal{J}_t,\mathbf{s}_t)$ which can be bounded using the Chernoff-Hoeffding bound and the properties of the UCB1 algorithm as follows,

$$P(\bar{C}(\mathcal{J}_t,\mathbf{s}_t)-\hat{C}(\mathcal{J}_t,\mathbf{s}_t)>\frac{\delta}{2})\leq\exp(-\frac{u(t)\delta^2}{2}), \quad (31)$$

where $u(t)$ is the total number of times the strategy $(\mathcal{J}_t,\mathbf{s}_t)$ has been used until time t. Since we use the UCB1 algorithm, in all the time instants t in which the arm $(\mathcal{J}_t,\mathbf{s}_t)$ has been chosen at least $\frac{8\log t}{\Delta_t^2}$ where ($\Delta_t=\bar{C}(\mathcal{J}^*,\mathbf{s}^*)-\bar{C}(\mathcal{J}_t,\mathbf{s}_t)$ is the sub-optimality gap of the strategy $(\mathcal{J}_t,\mathbf{s}_t)$), we have $P(\bar{C}(\mathcal{J}_t,\mathbf{s}_t)-\hat{C}(\mathcal{J}_t,\mathbf{s}_t)>\frac{\delta}{2})\leq\exp(-16\log t)=t^{-16}$, because $\Delta_t\leq\delta/2$ from the bound on the first term. This completes the proof of the Theorem.

APPENDIX E
JAMMING BANDITS WITH ARM ELIMINATION

Algorithm 3 is the modified JB algorithm that uses the UCB-Improved algorithm [34] instead of the UCB1 algorithm. Even in this algorithm, the regret has two terms, similar to the earlier regret

bounds presented in (10) and (11). While the bound in (10) continues to hold true even in this case, the bound in (11) changes due to the UCB-Improved algorithm. From [34], the regret bound for a M^2-armed bandit algorithm is given by $\mathcal{O}(\sqrt{M^2T}\frac{\log(M^2\log(M^2))}{\sqrt{\log(M^2)}})$. Thereby the overall regret is $TL\left(\frac{2}{M^2}\right)^{\frac{\alpha}{2}} + O\left(\sqrt{M^2T}\frac{\log(M^2\log(M^2))}{\sqrt{\log(M^2)}}\right)$. By using the doubling trick, the regret bounds for the overall time horizon n can be obtained.

APPENDIX F
JAMMING BANDITS AGAINST TIME-VARYING VICTIMS

For a given signaling scheme $\mathcal{J}$ used by the jammer, the average regret (i.e., with respect to the victim's strategies) can be bounded as below,

$$R_n = \Big[\sum_{t=1}^{n}\sum_{i=1}^{|\mathcal{P}|} p_i\Big(\mathbf{E}(C_t^i(\mathbf{s}^*)) - \mathbf{E}(C_t^i(\mathbf{s}_t))\Big)\Big] = \Big[\sum_{t=1}^{n}\sum_{i=1}^{|\mathcal{P}|} p_i\Big(\bar{C}_t^i(\mathbf{s}^*) - \bar{C}_t^i(\mathbf{s}_t)\Big)\Big]$$
$$= \Big[\sum_{t=1}^{n}\sum_{i=1}^{|\mathcal{P}|} p_i\Big(\{\bar{C}_t^i(\mathbf{s}^*) - \bar{C}_t^i(\mathbf{s}')\} + \{\bar{C}_t^i(\mathbf{s}') - \bar{C}_t^i(\mathbf{s}_t)\}\Big)\Big], \tag{32}$$

where $\bar{C}$ indicates the average cost function, $\mathbf{s}^*$ is the optimal jamming strategy across all stochastic strategies that can be used by the victim (since the jammer is not aware of the action taken by the user), $\mathbf{s}'$ is the strategy closest in Euclidean distance to the optimal strategy (as defined in the Section III) and $\mathbf{s}_t$ is the actual strategy chosen at time t. The first term can be bounded using the Hölder continuity properties of the cost function (see Appendix B where a similar analysis is done for the fixed user strategy case). The second term is bounded as follows. Since the jammer is not aware of the users' strategy, partition of the action space is done a priori and hence we use the same discretization M across all users' strategies. Thus define $\Delta = \sqrt{M^2\log(T)/T}$. For each action i taken by the user, define $\Delta_{ij} = \bar{C}(s') - \bar{C}(s^j)$ as the loss in rewards when a sub-optimal arm j is chosen. We split the arms into two sets, a) for those arms which satisfy $\Delta_{ij} < \Delta$, the near optimal arms and b) $\Delta_{ij} > \Delta$, the sub-optimal arms. For the first set, the maximum regret incurred against each user strategy is $T\Delta$ (over a time period T of the inner loop in Algorithm 1). For the second set, since we use the UCB1 algorithm, each sub-optimal arm is only chosen a maximum of $\frac{8\log(T)}{\Delta_{ij}^2} + 1 + \frac{\Pi^2}{3}$ times [26]. Thus we can bound the regret for the arms in the second set as $\Delta_{ij}\left(\frac{8\log(T)}{\Delta_{ij}^2} + 1 + \frac{\Pi^2}{3}\right)$. Overall the regret is upper bounded by $\sum_{i=1}^{|\mathcal{P}|} p_i\Big(TL\left(\frac{2}{M^2}\right)^{\alpha/2} + T\Delta + \sum_{j=1}^{M^2}\Delta_{ij}\left[\frac{8\log(T)}{\Delta_{ij}^2} + 1 + \frac{\Pi^2}{3}\right]\Big)$, which by using the relationship between

Δ and Δ_{ij} and the fact that $\Delta_{ij} \in [0,1]$ can further be upper bounded as

$$\left(TL\left(\frac{2}{M^2}\right)^{\alpha/2} + T\Delta + M^2\left[\frac{8\log(T)}{\Delta} + 1 + \frac{\Pi^2}{3}\right]\right) \approx \mathcal{O}(T^{\frac{\alpha+2}{2(\alpha+1)}}\log^{\frac{\alpha}{2(\alpha+1)}}T). \tag{33}$$

Since the jammer can choose from N_{mod} signaling schemes, the overall regret is the scaled (by N_{mod}) version of (33), which is similar to the regret bound in Theorem 2. Remember this is just a upper bound on the regret and that the actual regret depends on the strategy employed by the victim transmit-receive pair. Also notice that, all the regret and high confidence bounds derived in Section III can also be extended to this scenario as well.